arXiv:1602.06235v1 [stat.ML] 19 Feb 2016

A Mutual Contamination Analysis of Mixed Membership and Partial Label Models

Julian Katz-Samuels and Clayton Scott
{jkatzsam,clayscot}@umich.edu

Department of Electrical Engineering and Computer Science,
University of Michigan

February 22, 2016

Abstract

Many machine learning problems can be characterized by *mutual contamination models*. In these problems, one observes several random samples from different convex combinations of a set of unknown base distributions. It is of interest to *decontaminate* mutual contamination models, i.e., to recover the base distributions either exactly or up to a permutation. This paper considers the general setting where the base distributions are defined on arbitrary probability spaces. We examine the decontamination problem in two mutual contamination models that describe popular machine learning tasks: recovering the base distributions up to a permutation in a mixed membership model, and recovering the base distributions exactly in a partial label model for classification. We give necessary and sufficient conditions for identifiability of both mutual contamination models, algorithms for both problems in the infinite and finite sample cases, and introduce novel proof techniques based on affine geometry.

1 Introduction

In a mutual contamination model (Blanchard and Scott, 2014), there are L distributions $P_1, \ldots, P_L$ called *base distributions*. The learner observes M training random samples

$$X_1^i, \ldots, X_{n_i}^i \overset{i.i.d.}{\sim} \tilde{P}_i = \sum_{j=1}^{L} \pi_{i,j} P_j$$

$i = 1, \ldots, M$, where $\pi_{i,j} \geqslant 0$ and $\sum_j \pi_{i,j} = 1$. Here $\pi_{i,j}$ is the probability that an instance of the *contaminated distribution* $\tilde{P}_i$ is a realization of P_j. The $\pi_{i,j}$ and P_js are unknown and $\tilde{P}_i$ is observed through data. In this work, we avoid parametric models and assume that the sample space is arbitrary. The model can be stated concisely as

$$\tilde{\boldsymbol{P}} = \boldsymbol{\Pi} \boldsymbol{P} \tag{1}$$

where $\boldsymbol{P} = (P_1, \ldots, P_L)^T$, $\tilde{\boldsymbol{P}} = (\tilde{P}_1, \ldots, \tilde{P}_M)^T$, and $\boldsymbol{\Pi} = \left(\pi_{i,j}\right)$ is an $M \times L$ matrix (which we call the *mixing matrix*). The *decontamination* problem is to recover the base distributions either exactly or up to a permutation.

We study the decontamination problem in two mutual contamination models that describe popular machine learning tasks. First, in *mixed membership models*, the learner observes samples from $\tilde{P}_1, \ldots, \tilde{P}_M$ and the decontamination problem is to recover a permutation of $P_1, \ldots, P_L$. We refer to this problem as the decontamination or *demixing* problem interchangeably. Recently, mixed membership models have become a powerful modeling tool for data where data points are associated with multiple distributions. Applications have appeared in a wide range of fields including image processing (Li and Perona, 2005), population genetics (Pritchard et al., 2000), document analysis (Blei et al., 2003), and surveys (Berkman et al., 1989). There is potential value in developing a theory of mixed membership models for arbitrary sample spaces since in some applications (e.g., image processing and computer vision) it may be preferable to model features as more general random variables (Poczos et al., 2012). Such a theory will necessarily differ from previous work (reviewed below) which relies heavily on modeling the base and contaminated distributions as finite-dimensional probability vectors.

Consider the following application from natural language processing. In topic models, a form of mixed membership models, the P_is are distributions over words and the $\tilde{P}_i$s represent documents. As such, P_is are often treated as discrete random variables. However, it may be desirable to use word representations (mappings of words to $\mathbb{R}^d$) since word representations have contributed immensely to the success of systems on many tasks including parsing, entity recognition, and part-of-speech tagging (Luong et al., 2013). In such an approach, the P_is are continuous distributions over $\mathbb{R}^d$. Whereas most topic modeling algorithms could not operate on such a useful representation, our results and algorithms do apply.

A partial label model is an alternative setting for multiclass classification with L classes.[1] In a partial label model, each data point is labeled with a *partial label* $S \subset \{1, \ldots, L\}$; the true label is in S, but is it not known which label is the true one. If S_i is a partial label, then we view patterns with label S_i as distributed according to $\tilde{P}_i = \sum_{j \in S_i} \pi_{i,j} P_j$. In the mutual contamination model setting, the learner has access to the *partial label matrix* $\boldsymbol{S} = (\mathbf{1}_{\{\boldsymbol{\Pi}_{i,j} > 0\}})$, an $M \times L$ matrix. The decontamination problem is to recover $(P_1, \ldots, P_L)^T$ *exactly* (not just up to a permutation) from $\boldsymbol{S}$ and $\tilde{\boldsymbol{P}}$ (or its empirical version).

Most work has approached the partial label problem by trying to minimize the partial label error (Jin and Ghahramani, 2002; Nyugen and Caruana, 2008; Cour et al., 2011; Liu and Dietterich, 2012). The partial label error is the probability that a given classifier assigns a label to a training instance that is not contained in the partial label associated with the training instance. To our knowledge our work is the first to consider recovery of the base distributions P_j. Once the base distributions are known, it would then be possible to design a classifier according to any criterion, such as probability of error.

We make the following contributions: (i) We give necessary and sufficient conditions on $\boldsymbol{P}$, $\boldsymbol{\Pi}$, and $\boldsymbol{S}$ for identifiability of the mixed membership and partial label models under mild assumptions on the decontamination procedure. (ii) We introduce novel algorithms for the demixing problem and partial label classification in the infinite and finite sample settings. These algorithms are nonparametric in the sense that they do not model P_i as a probability vector or other parametric model. (iii) We develop novel estimators for distributions obtained by iteratively applying the κ^* operator (defined below). (iv) We introduce novel proof techniques based on affine geometry.

1.1 Notation

Let $\mathbb{Z}^+$ denote the positive integers. For $n \in \mathbb{Z}^+$, let $[n] = \{1, \ldots, n\}$. If $\boldsymbol{x} \in \mathbb{R}^K$, let x_i denote the ith entry of $\boldsymbol{x}$. In contrast, if $\boldsymbol{x}_j \in \mathbb{R}^K$, then $x_{j,i}$ denotes the ith entry of $\boldsymbol{x}_j$. Let $\boldsymbol{e}_i$ denote the length L vector with 1 in the ith position and zeros elsewhere. Let $\boldsymbol{\pi}_i \in \Delta_L \subset \mathbb{R}^L$ be the transpose of the

[1] The partial label problem has also been referred to as the "superset learning problem" or the "multiple label problem" (Liu and Dietterich, 2014).

ith row of $\mathbf{\Pi}$ where Δ_L denotes the $(L-1)$-dimensional simplex, i.e., $\Delta_L = \{\boldsymbol{\mu} = (\mu_1, \ldots, \mu_L)^T \in \mathbb{R}^L \mid \sum_{i=1}^L \mu_i = 1 \text{ and } \forall i : \mu_i \geqslant 0\}$. Let Δ_L^M denote the product of M $(L-1)$-dimensional simplices, the space of $M \times L$ stochastic matrices. Let $\mathcal{P}$ denote the space of probability distributions on a measurable space $(\mathcal{X}, \mathcal{C})$. Let $\operatorname{supp}(F)$ denote the support of a distribution F on a Borel space. Let $\boldsymbol{B}^{M \times L}$ denote the space of binary matrices of dimension $M \times L$. Let A be a set. Let $\operatorname{aff} A$ denote the affine hull of A, i.e., $\operatorname{aff} A = \{\sum_{i=1}^K \theta_i \boldsymbol{x}_i | \boldsymbol{x}_1, \ldots, \boldsymbol{x}_K \in A, \sum_{i=1}^K \theta_i = 1\}$.

2 Related Work

Mutual contamination models were employed by Scott et al. (2013) and Blanchard and Scott (2014) to study classification with label noise. Scott et al. (2013) assume that $M = L = 2$ and show how to recover P_1 and P_2 exactly (not just up to a permutation). Blanchard and Scott (2014) assume that $M = L \geqslant 2$ and recover $P_1, \ldots, P_L$ exactly. The demixing problem in the current paper differs from these problems in that *(i)* the number of contaminated distributions is not necessarily equal to the number of base distributions (i.e., we allow $M \neq L$) and, more significantly, *(ii)* the demixing problem only requires recovering $P_1, \ldots, P_L$ up to a permutation. The partial label model differs from these problems in that it assumes access to the partial label matrix $\boldsymbol{S}$. We use some ideas from (Blanchard and Scott, 2014), but make significant extensions. These differences lead us to make different (substantially weaker) assumptions on $\mathbf{\Pi}$ to achieve identifiability, although we use the same assumption on $\boldsymbol{P}$, namely that $P_1, \ldots, P_L$ are jointly irreducible (defined below).

2.1 The Demixing Problem and Topic Models

Our demixing problem may also be viewed as topic modeling on general domains. In topic modeling, the base distributions P_i correspond to topics and the contaminated distributions $\tilde{P}_i$ to documents, which are regarded as mixtures of topics. In most cases, the P_is are assumed to have a finite sample space. A variety of approaches have been proposed for topic modeling. The most common approach assumes a generative model for a corpus of documents and determines the maximum likelihood fit of the model given data. However, because maximum likelihood is NP-hard, these approaches must rely on heuristics that can get stuck in local minima (Arora et al., 2012).

Recently, a trend towards algorithms for topic modeling with provable guarantees has emerged. These methods rely on the separability assumption **(SEP)** and its variants. According to **(SEP)**, $P_1, \ldots, P_L$ are distributions on a finite sample space and for every $i \in \{1, \ldots, L\}$, there exists a word $x \in \operatorname{supp}(P_i)$ such that $x \notin \cup_{j \neq i} \operatorname{supp}(P_j)$. Our requirement that $P_1, \ldots, P_L$ are jointly irreducible is a natural generalization of separability of $P_1, \ldots, P_L$, as we will argue below. Specifically, if $P_1, \ldots, P_L$ have discrete sample spaces, separability and joint irreducibility coincide; however, if $P_1, \ldots, P_L$ are continuous, under joint irreducibility, $P_1, \ldots, P_L$ can have the same support.

A key ingredient in these algorithms is to use the assumption of a finite sample space to view the distributions as probability vectors in Euclidean space; this leads to approaches based on non-negative matrix factorization (NMF), linear programs, and random projections (Donoho and Stodden, 2003; Arora et al., 2012, 2013; Ding et al., 2013, 2014; Recht et al., 2012). However, more general distributions cannot be viewed as finite-dimensional vectors. Therefore, topic modeling on general domains requires new techniques. Our work seeks to provide such techniques.

2.2 Partial Label Model

The partial label model has had two main formulations in previous work. In **(PL-1)**, instances from each class are drawn independently and the partial label for each instance is drawn independently from a set-valued distribution. In **(PL-2)**, training data are in the form of bags where each bag is a set of instances and the bag has a set of labels. Each instance belongs to a single class, and the set of labels associated with the bag is given by the union of the labels of the instances in the bag (Liu and Dietterich, 2014). Our framework is similar to **(PL-2)**, although it does not assume a joint distribution on the features of instances and the partial labels.

Most algorithms approach classification in a partial label model by picking a classifier that minimizes the partial label error (Jin and Ghahramani, 2002; Nyugen and Caruana, 2008; Cour et al., 2011; Liu and Dietterich, 2012). Liu and Dietterich (2014) develop learnability results in the realizable case for algorithms that use this approach. One of the key concepts that they take from Cour et al. (2011) is *the ambiguity degree*. It bounds the probability that a specific incorrect label appears in the partial label of an instance. Our approach has the advantage that it makes no assumption regarding realizability, which essentially means that P_i have disjoint supports.

Our paper makes three main contributions to the literature. First, we address the fundamental question of identifiability of a partial label model. Second, we provide nonparametric algorithms for the infinite and finite sample settings. Finally, we introduce a novel approach based on affine geometry that differs significantly from approaches that minimize the partial label error.

3 Necessary Conditions

We begin by developing necessary conditions for identifiability of a mixed membership model and, then, discuss necessary conditions for the identifiability of a partial label model. A mixed membership model is identifiable if, given $\tilde{\boldsymbol{P}}$, the pair $(\boldsymbol{\Pi}, \boldsymbol{P})$ that solves (1) is uniquely determined. In general, this pair is not unique. For example, consider the case where $L = M$, $(\boldsymbol{\Pi}, \boldsymbol{P})$ solves (1), and $\boldsymbol{\Pi}$ is not a permutation matrix. Then, another solution is $\tilde{\boldsymbol{P}} = \boldsymbol{I}\tilde{\boldsymbol{P}}$.

Furthermore, there are infinitely many non-trivial solutions in the realistic scenario where there is some $\tilde{P}_i$ in the interior of $\text{conv}(P_1, \ldots, P_L)$ and at least two of the $\tilde{P}_j$s are distinct. We can construct such solutions as follows. Without loss of generality, suppose that $i = 1$ and $\tilde{P}_1 \neq \tilde{P}_2$. Then, since $\tilde{P}_1$ is in the interior of $\text{conv}(P_1, \ldots, P_L)$, there is some $\delta > 0$ such that for any $\alpha \in (1, 1+\delta)$, $Q_\alpha = \alpha\tilde{P}_1 + (1-\alpha)\tilde{P}_2$ is a distribution. Then, $\text{conv}(\tilde{P}_1, \ldots, \tilde{P}_L) \subseteq \text{conv}(Q_\alpha, \tilde{P}_2, \ldots, \tilde{P}_L)$ and, consequently, there is some $\boldsymbol{\Pi}' \in \Delta_L^L$ such that $(\boldsymbol{\Pi}', (Q_\alpha, \tilde{P}_2, \ldots, \tilde{P}_L)^T)$ solves (1). Clearly, by varying α, there are infinitely many solutions to (1). Moreover, we can replace $\tilde{P}_2$ in the above argument with any distribution in the convex hull of $\{\tilde{P}_j : j \neq 1\}$ that is not equal to $\tilde{P}_1$.

In light of the above, it is natural to impose conditions on the decontamination procedure so as to eliminate trivial and other simple solutions to (1). We now formalize the notion of a decontamination operator and present two conditions that we believe it should satisfy. Let ϕ denote a *decontamination operator*: a function from $\mathcal{P}^M$ to $\Delta_L^M \times \mathcal{P}^L$ such that $\phi(\tilde{\boldsymbol{P}})$ returns $(\boldsymbol{\Pi}, \boldsymbol{P})$ that solves (1). Let $\phi_1(\tilde{\boldsymbol{P}})$ return $\boldsymbol{\Pi}$ and $\phi_2(\tilde{\boldsymbol{P}})$ return $\boldsymbol{P}$.

Definition 1. *A decontamination operator ϕ satisfies Maximality **(M)** iff $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$ implies that any pair $(\boldsymbol{\Pi}', \boldsymbol{P}')$ with $\boldsymbol{\Pi}' \in \Delta_L^M$ and $\boldsymbol{P}' = (P_1', \ldots, P_L')^T \in \mathcal{P}^L$ that solves (1) satisfies $\{P_1', \ldots, P_L'\} \subseteq \text{conv}(P_1, \ldots, P_L)$.*

In words, **(M)** states that $\boldsymbol{P} = (P_1, \ldots, P_L)^T$ is a maximal collection of base distributions in the sense that it is not possible to move any of the P_is outside of $\text{conv}(P_1, \ldots, P_L)$ and represent $\tilde{\boldsymbol{P}}$.

Definition 2. *A decontamination operator ϕ satisfies Linearity **(L)** iff $\phi_2(\tilde{\boldsymbol{P}}) = \boldsymbol{P}$ implies $\{P_1, \ldots, P_L\} \subseteq \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M)$.*

We believe that **(L)** is a reasonable requirement since it holds in the common situation in which there exist $\boldsymbol{\pi}_{i_1}, \ldots, \boldsymbol{\pi}_{i_L}$ that are linearly independent. Then for $I = \{i_1, \ldots, i_L\}$, we can write $\tilde{\boldsymbol{P}}_I = \boldsymbol{\Pi}_I \boldsymbol{P}$ where $\boldsymbol{\Pi}_I$ is the submatrix of $\boldsymbol{\Pi}$ containing only the rows indexed by I and $\tilde{\boldsymbol{P}}_I$ is similarly defined. Then, $\boldsymbol{\Pi}_I$ is invertible and $\boldsymbol{P} = \boldsymbol{\Pi}_I^{-1} \tilde{\boldsymbol{P}}$.

To formulate a necessary condition, we introduce another definition. For distributions G and H, we say that G is *irreducible* with respect to H if it is not possible to write $G = \gamma H + (1 - \gamma) F$ where F is a distribution and $0 < \gamma \leqslant 1$. This condition was used previously in the study of classification with label noise (Blanchard et al., 2010; Scott et al., 2013); here, we show that it arises in the context of a necessary condition for demixing mixed membership models.

Theorem 1. *Let ϕ denote a decontamination operator and $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. If ϕ satisfies **(M)**, then*

(A) *$\forall i$, P_i is irreducible with respect to every distribution in* $\text{conv}(\{P_j : j \neq i\})$.

*If ϕ satisfies **(L)**, then*

(B) $\text{rank}(\boldsymbol{\Pi}) \geqslant \dim \text{span}(P_1, \ldots, P_L)$.

Sketch. **(A)** Proof by contraposition. Let ϕ be a demixing operator and $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. Suppose that there is some P_i and $Q \in \text{conv}(\{P_j : j \neq i\})$ such that P_i is not irreducible with respect to Q. Then, there is some distribution G and $\gamma \in (0, 1]$ such that $P_i = \gamma Q + (1 - \gamma) G$. First, we show that if $\gamma = 1$ or $G \in \text{conv}(P_1, \ldots, P_L)$, then $P_i \in \text{conv}(\{P_k : k \neq i\})$ in which case $\tilde{P}_1, \ldots, \tilde{P}_M \in \text{conv}(\{P_j : j \neq i\} \cup \{R\})$ for any distribution $R \notin \text{conv}(P_1, \ldots, P_L)$. Thus, ϕ violates **(M)**. Second, we show that if $\gamma \in (0, 1)$ and $G \notin \text{conv}(P_1, \ldots, P_L)$, then $\text{conv}(P_1, \ldots, P_L) \subseteq \text{conv}(\{P_k : k \neq i\} \cup \{G\})$, from which it follows that ϕ violates **(M)**.

(B) Let $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. Using the hypothesis that ϕ satisfies **(L)** and the relation $\tilde{\boldsymbol{P}} = \boldsymbol{\Pi}\boldsymbol{P}$, we show that $\dim \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) = \dim \text{span}(P_1, \ldots, P_L)$. Then, the result follows from the fact that $\tilde{P}_1, \ldots, \tilde{P}_M \in \text{range}(\boldsymbol{\Pi})$.

□

In sum, we have obtained necessary conditions on the base distributions and the mixing matrix for identifiability of a mixed membership model. In particular, Theorem 1 implies that if $P_1, \ldots, P_L$ are linearly independent, then there must be at least as many contaminated distributions as base distributions, i.e., $M \geqslant L$. We also remark that **(A)** and **(B)** apply to Blanchard and Scott (2014) (in which no necessary conditions were given) since in that paper, the goal is to identify $P_1, \ldots, P_L$ exactly. Further, note that **(A)** appears as a sufficient condition in Sanderson and Scott (2014).

Now, we turn to the partial label model. A partial label model is identifiable if given $\tilde{\boldsymbol{P}}$ and $\boldsymbol{S}$, the pair $(\boldsymbol{\Pi}, \boldsymbol{P})$ that is consistent with $\boldsymbol{S}$ and solves $\tilde{\boldsymbol{P}} = \boldsymbol{\Pi}\boldsymbol{P}$ is unique. A decontamination operator for the partial label model ψ is a function from $\mathcal{P}^M \times \boldsymbol{B}^{M \times L}$ to $\Delta_L^M \times \mathcal{P}^L$. This difference necessitates slight modifications to our notions of Maximality and Linearity, which we refer to as **(M′)** and **(L′)**, respectively. Let $\psi(\tilde{\boldsymbol{P}}, \boldsymbol{S}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. If ψ satisfies **(M′)**, then $(\boldsymbol{\Pi}, \boldsymbol{P})$ need not satisfy **(A)**. However, if ψ satisfies **(B)**, then $(\boldsymbol{\Pi}, \boldsymbol{P})$ satisfies **(B)**. See the Appendix for details. Identifiability of the partial label model implies a condition **(C)** on the partial label matrix.

Proposition 1. *Let ψ denote a decontamination operator for the partial label model and $\psi(\tilde{\boldsymbol{P}}, \boldsymbol{S}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. If $(\tilde{\boldsymbol{P}}, \boldsymbol{S})$ is identifiable, then*

(C) *$\boldsymbol{S}$ does not contain a pair of identical columns.*

Sketch. Let $(\tilde{\boldsymbol{P}}, \boldsymbol{S})$ such that $\boldsymbol{S}$ does not satisfy **(C)**. Without loss of generality, suppose that the $\boldsymbol{S}_{:,1}$ and $\boldsymbol{S}_{:,2}$ are identical. Let $(\boldsymbol{\Pi}, \boldsymbol{P})$ satisfy (1). Interchange the first two columns of $\boldsymbol{\Pi}$ to obtain $\boldsymbol{\Pi}'$ and the first two rows of $\boldsymbol{P}$ to obtain $\boldsymbol{P}'$. Then, $(\boldsymbol{\Pi}', \boldsymbol{P}')$ solves (1), establishing that $(\tilde{\boldsymbol{P}}, \boldsymbol{S})$ is not identifiable. □

4 Sufficient Conditions

Now, we state the sufficient conditions for demixing a mixed membership model and decontaminating a partial label model. Consider the following property from Blanchard and Scott (2014).

Definition 3. *The distributions $\{P_i\}_{1 \leqslant i \leqslant L}$ are* jointly irreducible *iff the following equivalent conditions hold*

(a) It is not possible to write

$$\sum_{i \in I} \epsilon_i P_i = \alpha \sum_{i \notin I} \epsilon_i P_i + (1 - \alpha) H$$

where $I \subset [L]$ such that $1 \leqslant |I| < L$, ϵ_i are such that $\epsilon_i \geqslant 0$ and $\sum_{i \in I} \epsilon_i = \sum_{i \notin I} \epsilon_i = 1$, $\alpha \in (0, 1]$ and H is a distribution;

(b) $\sum_{i=1}^{L} \gamma_i P_i$ is a distribution implies that $\gamma_i \geqslant 0 \, \forall i$.

Conditions *(a)* and *(b)* give two ways to think about joint irreducibility. Condition *(a)* says that every convex combination of a subset of the P_is is irreducible with respect to every convex combination of the other P_is. Condition *(b)* says that if a distribution is in the span of $P_1, \ldots, P_L$, it is in their convex hull. Joint irreducibility holds when each P_i has a region of positive probability that does not belong to the support of any of the other P_is; thus, separability of P_is entails joint irreducibility of $P_1, \ldots, P_L$. However, the converse is not true: the P_is can have the same support and still be jointly irreducible (e.g., P_is Gaussian with a common variance and distinct means (Scott et al., 2013)).

Our two sufficient conditions for the mixed membership model are:

(A′) $P_1, \ldots, P_L$ are jointly irreducible.

(B′) $\boldsymbol{\Pi}$ has full rank.

These conditions are consistent with requiring that a decontamination operator satisfy **(M)** and **(L)**, as indicated by the following Proposition.[2]

Proposition 2. *If a decontamination operator ϕ is such that $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$ only if $(\boldsymbol{\Pi}, \boldsymbol{P})$ satisfy **(A′)** and **(B′)**, then ϕ satisfies **(M)** and **(L)**.*

By comparing **(A)** with **(A′)** and **(B)** with **(B′)**, we see that the proposed sufficient conditions are not much stronger than **(M)** and **(L)** require. Since joint irreducibility of $P_1, \ldots, P_L$ entails their linear independence by Lemma B.1 of Blanchard and Scott (2014), under **(A′)**, **(B)** and **(B′)** are the same. **(A′)** differs from **(A)** in that it requires that a slightly larger set of distributions are irreducible with respect to convex combinations of the remaining distributions. Specifically, under **(A′)**, *every convex combination* of a subset of the P_is is

[2] We give an analogous result for the partial label model in the Appendix.

irreducible with respect to every convex combination of the other P_is whereas **(A)** only requires that every P_i be irreducible with respect to every convex combination of the other P_is.

Our sufficient conditions for the partial label model include **(A′)**, **(B′)**, as well as **(C)**. Note that our sufficient and necessary conditions on $\boldsymbol{S}$ are identical.

A couple of points are in order regarding how **(A′)** and **(B′)** relate to the sufficient conditions in Blanchard and Scott (2014). Blanchard and Scott (2014) also assume **(A′)**. However, the so-called *recoverability* assumption on $\boldsymbol{\Pi}$ in Blanchard and Scott (2014) is substantially stronger than **(B′)**. Recoverability says that $\boldsymbol{\Pi}^{-1}$ has a very specific structure: positive diagonal entries and nonpositive off-diagonal entries. To demix a mixed membership model, we are able to weaken the assumption on $\boldsymbol{\Pi}$ because the goal is to recover the base distributions only up to a permutation (as opposed to a specific permutation of the base distributions). Recall that in the classification problem in a partial label model, the goal is to recover $P_1, \ldots, P_L$ exactly. In comparison to Blanchard and Scott (2014), we are able to weaken the assumption on $\boldsymbol{\Pi}$ because the partial label matrix $\boldsymbol{S}$ gives a considerable amount of useful information.

5 Population Case

In this section, to establish that the above conditions are indeed sufficient, we show that under these assumptions, mixed membership models and partial label models can be decontaminated in the population case. Henceforth, we assume that $P_1, \ldots, P_L$ satisfy are jointly irreducible, $\boldsymbol{\Pi}$ has full rank, and $M = L$. We give extensions to the non-square case in the Appendix. We begin by reviewing the necessary background.

5.1 Background

This paper relies on two important quantities from Blanchard et al. (2010) and Blanchard and Scott (2014), which we now review.

Proposition 3. *Given probability distributions F_0, F_1 on a measurable space $(\mathcal{X}, \mathcal{C})$, define*

$$\kappa^*(F_0 \,|\, F_1) = \max\{\kappa \in [0,1] | \, \exists \text{ a distribution } G \text{ s.t. } F_0 = (1-\kappa)G + \kappa F_1\};$$

If $F_0 \neq F_1$, then $\kappa^*(F_0 \,|\, F_1) < 1$ and the above supremum is attained for a unique distribution G (which we refer to as the residue *of F_0 wrt. F_1). Furthermore, the following equivalent characterization holds:*

$$\kappa^*(F_0 \,|\, F_1) = \inf_{C \in \mathcal{C}, F_1(C) > 0} \frac{F_0(C)}{F_1(C)}.$$

$\kappa^*(F_0 \,|\, F_1)$ can be thought of as the maximum possible proportion of F_1 in F_0. Note that $\kappa^*(F_0 \,|\, F_1) = 0$ iff F_0 is irreducible wrt F_1. We can think of $1 - \kappa^*(F_0 \,|\, F_1)$ as a statistical distance since it it non-negative and equal to zero if and only if $F_0 = F_1$. We refer to this quantity as the two-sample κ^* operator. To obtain the residue of F_0 wrt F_1, one computes Residue$(F_0 \,|\, F_1)$ (see Algorithm 4); this is well-defined under Proposition 3 when $F_0 \neq F_1$.

We now turn to the multi-sample generalization of κ^* defined in Blanchard and Scott (2014), which we call the multi-sample κ^* operator.

Definition 4. *Given distributions $F_0, \ldots, F_K$, define $\kappa^*(F_0 \,|\, F_1, \ldots, F_K) =$*

$$\max\Big(\sum_{i=1}^K \nu_i : \nu_i \geqslant 0, \sum_{i=1}^K \nu_i \leqslant 1, \exists \text{ distribution } G \text{ s.t. } F_0 = (1 - \sum_{i=1}^K \nu_i)G + \sum_{i=1}^K \nu_i F_i\Big)$$

Under our sufficient conditions, there exists some G attaining the above maximum. However, the G is not necessarily unique. Any G attaining the maximum is called a multi-sample residue of F_0 wrt $F_1, \ldots, F_K$. The algorithm Residue($F_0 \mid \{F_1, \ldots, F_K\}$) returns one of these G (see Algorithm 5).

In previous work that assumes P_i are probability vectors, distributions are compared using l_p distances. In our setting of general probability spaces, we use κ^* to compare different distributions.

5.2 Mixture Proportions

If $\boldsymbol{\eta} \in \mathbb{R}^L$ and $Q = \boldsymbol{\eta}^T \boldsymbol{P}$, we say that $\boldsymbol{\eta}$ is the *mixture proportion* of Q. Since by Lemma B.1 of Blanchard and Scott (2014), joint irreducibility of $P_1, \ldots, P_L$ implies linear independence of $P_1, \ldots, P_L$, mixture proportions are well-defined.

An important feature of our solution strategy involves determining where the mixture proportions of distributions are in the simplex Δ_L. To make this precise, we introduce the following definitions. If $i \in [L]$, we say that $\text{conv}(\{\boldsymbol{e}_j : j \neq i\})$ is a *face* of the simplex Δ_L; if $A \subset [L]$ and $|A| = k$, we also say that $\text{conv}(\{\boldsymbol{e}_j : j \in A\})$ is a *k-face* of Δ_L. If $\boldsymbol{\eta} \in \mathbb{R}^L$, Q is a distribution, and $Q = \boldsymbol{\eta}^T \boldsymbol{P}$, we say that $\mathcal{S}(\boldsymbol{\eta}) = \{j : \eta_j > 0\}$ is the *support set* of $\boldsymbol{\eta}$ or the support set of Q. Note that in this setting, $\mathcal{S}(\boldsymbol{\eta})$ consists of the indices of all the nonzero entries in the mixture proportion $\boldsymbol{\eta}$ by definition of joint irreducibility. Finally, for $\boldsymbol{\eta}_i \in \Delta_L$, and $Q_i = \boldsymbol{\eta}_i^T \boldsymbol{P}$ for $i = 1, 2$, we say that the distributions Q_1 and Q_2 (or the mixture proportions $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_2$) are *on the same face* of the simplex Δ_L if there exists $j \in [L]$ such that $\boldsymbol{\eta}_1, \boldsymbol{\eta}_2 \in \text{conv}(\{\boldsymbol{e}_k : k \neq j\})$.

The heart of our approach is that under joint irreducibility, one can interchange distributions $Q_1, \ldots, Q_K$ and their mixture proportions $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_K$, as indicated by the following Proposition.

Proposition 4. *Let $Q_i = \boldsymbol{\eta}_i^T \boldsymbol{P}$ for $i \in [L]$ and $\boldsymbol{\eta}_i \in \Delta_L$. If $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_L$ are linearly independent and $P_1, \ldots, P_L$ are jointly irreducible, then for any $i \in [L]$ and $A \subseteq [L] \backslash \{i\}$, $\kappa^*(Q_i \mid \{Q_j : j \in A\}) = \kappa^*(\boldsymbol{\eta}_i \mid \{\boldsymbol{\eta}_j : j \in A\}) < 1$. Further, $\boldsymbol{\gamma} \in \Delta_L$ is a residue of $\boldsymbol{\eta}_i$ wrt $\{\boldsymbol{\eta}_j : j \in A\}$ if and only if $G = \boldsymbol{\gamma}^T \boldsymbol{P}$ is a residue of Q_i wrt $\{Q_j : j \in A\}$.*

In words, this proposition says that the optimization problem given by $\kappa^*(Q_i \mid \{Q_j : j \in A\})$ is equivalent to the optimization problem given by $\kappa^*(\boldsymbol{\eta}_i \mid \{\boldsymbol{\eta}_j : j \in A\})$. Thus, joint irreducibility of $P_1, \ldots, P_L$ and linear independence of the mixture proportions ensure that we can reduce the decontamination problem of a mutual contamination model to the geometric problem of recovering the vertices of a simplex by applying κ^* to points (i.e., the mixture proportions) in the simplex. This makes the figures below valid for general distributions (see Figures 1, 2, and 3).

5.3 The Demix Algorithm

In essence, the Demix algorithm (see Algorithm 1) computes the inverse of $\boldsymbol{\Pi}$ in a sequential fashion using κ^* (recall we are assuming $M = L$). It is a recursive algorithm. Let $S_1, \ldots, S_K$ denote K contaminated distributions. In the base case, the algorithm takes as its input two contaminated distributions S_1 and S_2. It returns Residue($S_1 \mid S_2$) and Residue($S_2 \mid S_1$), which are a permutation of the two base distributions (see Figure 1). When $K > 2$, Demix finds $K - 1$ distributions $R_2, \ldots, R_K$ on the same $(K-1)$-face using a subroutine FaceTest (see Algorithm 6) and recursively applies Demix to $R_2, \ldots, R_K$ to obtain distributions $Q_1, \ldots, Q_{K-1}$. $Q_1, \ldots, Q_{K-1}$ are a permutation of $K - 1$ of the base distributions. Subsequently, the algorithm computes a sequence of residues using $Q_1, \ldots, Q_{K-1}$ and $\frac{1}{K} \sum_{i=1}^K S_i$ to obtain Q_K, the remaining base distribution (see Figure 2 for an execution of the algorithm). A number of remarks are in order. First, although we compute the residue of $\frac{1}{n} S_i + \frac{n-1}{n} Q$ wrt S_1 for each $i \neq 1$, there is nothing special about the

Algorithm 1 Demix$(S_1, \ldots, S_K)$

Input: $S_1, \ldots, S_K$ are distributions

if $K = 2$ **then**
  $Q_1 \longleftarrow$ Residue$(S_1 \mid S_2)$
  $Q_2 \longleftarrow$ Residue$(S_2 \mid S_1)$
  return $(Q_1, Q_2)^T$
else
  $Q \longleftarrow$ unif distributed element in conv$(S_2, \ldots, S_K)$
  $n \longleftarrow 1$
  $T \longleftarrow 0$
  while $T == 0$ **do**
    $n \longleftarrow n + 1$
    for $i = 2, \ldots, K$ **do**
      $R_i \longleftarrow$ Residue$(\frac{1}{n} S_i + \frac{n-1}{n} Q \mid S_1)$
    end for
    $T \longleftarrow$ FaceTest$(R_2, \ldots, R_K)$
  end while
  $(Q_1, \ldots, Q_{K-1})^T \longleftarrow$ Demix$(R_2, \ldots, R_K)$
  $Q_K \longleftarrow \frac{1}{K} \sum_{i=1}^{K} S_i$
  for $i = 1, \ldots, K-1$ **do**
    $Q_K \longleftarrow$ Residue$(Q_K \mid Q_i)$
  end for
  return $(Q_1, \ldots, Q_K)^T$
end if

distribution S_1. We could replace S_1 with any S_j where $j \in [K]$, provided that we adjust the rest of the algorithm accordingly. Second, we can replace the sequence $\{\frac{n-1}{n}\}_{n=1}^{\infty}$ with any sequence $\alpha_n \nearrow 1$. Third, Q is sampled randomly to ensure w.p. 1 it does not lie on a union of $k-2$ dimensional affine subspaces (a zero measure set) on which the algorithm fails. Specifically, under the assumptions of Theorem 2, w.p. 1, the residue of Q wrt S_1 lies on the interior of one of the $(K-1)$-faces, so if we pick distributions close enough to Q, their residues wrt S_1 lie on the same face.

Theorem 2. *Let $P_1, \ldots, P_L$ be jointly irreducible, $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_L \in \Delta_L$ be linearly independent, and $\tilde{P}_i = \boldsymbol{\pi}_i^T \boldsymbol{P}$ for $i \in [L]$. Then, with probability 1, Demix$(\tilde{P}_1, \ldots, \tilde{P}_L)$ terminates and returns a permutation of $(P_1, \ldots, P_L)^T$.*

5.4 The Partial Label Algorithm

The PartialLabel algorithm (see Algorithm 2) proceeds by iteratively finding sets of candidate distributions $(W_1, \ldots, W_L)^T$ for increasing values of k. Given each $(W_1, \ldots, W_L)^T$, it runs an algorithm VertexTest (see Algorithm 9) that uses $\tilde{P}_1, \ldots, \tilde{P}_L$ and the partial label matrix $\boldsymbol{S}$ to determine whether $(W_1, \ldots, W_L)^T$ is a permutation of the base distributions $(P_1, \ldots, P_L)^T$. If $(W_1, \ldots, W_L)^T$ is a permutation of $(P_1, \ldots, P_L)^T$, VertexTest constructs the corresponding permutation matrix for relating these distributions. If not, it returns a value indicating that the PartialLabel algorithm increment k to find another candidate set of distributions.

Algorithm 2 PartialLabel($\boldsymbol{S}, (\tilde{P}_1, \ldots, \tilde{P}_L)^T$)

for $i = 1, \ldots, L$ **do**
  $Q_i \longleftarrow$ uniformly random distribution in conv$(\tilde{P}_1, \ldots, \tilde{P}_L)$
  $W_i \longleftarrow Q_i$
end for
$k = 2$
FoundVertices $\longleftarrow 0$
while FoundVertices $== 0$ **do**
  for $i = 1, \ldots, L$ **do**
    $\bar{Q}_i \longleftarrow \frac{1}{L-1}[\sum_{j>i} Q_j + \sum_{j<i} W_j]$
    $W_i \longleftarrow$ Residue$(\frac{1}{k}Q_i + (1 - \frac{1}{k})\bar{Q}_i \,|\, \{Q_j\}_{j>i} \cup \{W_j\}_{j<i})$
  end for
  $k = k + 1$
  (FoundVertices, $\boldsymbol{C}$) $\longleftarrow$ VertexTest$(\boldsymbol{S}, \tilde{P}_1, \ldots, \tilde{P}_L, W_1, \ldots, W_L)$
end while
return $\boldsymbol{C}(W_1, \ldots, W_L)^T$

Theorem 3. *Suppose that $P_1, \ldots, P_L$ satisfy **(A′)**, $\boldsymbol{\Pi}$ satisfies **(B′)**, and $\boldsymbol{S}$ satisfies **(C)**. If $(R_1, \ldots, R_L)^T \longleftarrow$ PartialLabel$(\boldsymbol{S}, (\tilde{P}_1, \ldots, \tilde{P}_L)^T)$, then $R_i = P_i$ for all $i \in [L]$.*

Sketch. We adopt the notation from the description of the algorithm with the exception that we make explicit the dependence on k by writing $W_i^{(k)}$ instead of W_i and $\bar{Q}_i^{(k)}$ instead of $\bar{Q}_i$. In this proof sketch, we only show that there is a K such that for all $k \geqslant K$, $(S_1^{(k)}, \ldots, S_L^{(k)})^T$ is a permutation of $(P_1, \ldots, P_L)^T$.

Let $Q_i = \boldsymbol{\tau}_i^T \boldsymbol{P}$, $\bar{Q}_i^{(k)} = \boldsymbol{\tau}^{(k)^T} \boldsymbol{P}$, and $W_i^{(k)} = \boldsymbol{\gamma}_i^{(k)^T} \boldsymbol{P}$. We prove the claim inductively. First, we prove the base case: there is large enough k such that $W_1^{(k)} = P_i$ for some $i \in [L]$. By Lemma 7, $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_L$ are linearly independent. Therefore, aff$(\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L)$ gives a hyperplane with an associated open halfspace $\boldsymbol{H}$ that contains $\boldsymbol{\tau}_1$ and at least one $\boldsymbol{e}_j$. We can pick k large enough such that for all $\boldsymbol{e}_j \in \boldsymbol{H}$, $\lambda_k \equiv \frac{1}{k}\boldsymbol{\tau}_1 + \frac{k-1}{k}\boldsymbol{\tau}_1^{(k)} \in$ conv$(\boldsymbol{e}_j, \boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L)$. Then, for all $\boldsymbol{e}_j \in \boldsymbol{H}$, there exists $\kappa_j > 0$ such that

$$\boldsymbol{\lambda}_k = \kappa_j \boldsymbol{e}_j + (1 - \kappa_j)\tilde{\tau}_j$$

for some $\tilde{\tau}_j \in$ conv$(\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L)$. Suppose that there is $\boldsymbol{e}_i, \boldsymbol{e}_j \in \boldsymbol{H}$ such that $\kappa_i = \kappa_j$. Then, it can be shown that $\boldsymbol{e}_j - \boldsymbol{e}_i, \boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L$ are linearly dependent, which cannot happen by the randomness of $\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L$ (Lemma 7). Therefore, there is a unique minimum κ_j that satisfies the above relation. By Lemma 5, $\boldsymbol{e}_j$ is the residue of $\boldsymbol{\lambda}_k$ wrt $\{\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L\}$. By Proposition 4, it follows that $W_1^{(k)}$ is one of the base distributions. □

The VertexTest algorithm proceeds as follows on a vector of candidate distributions $(Q_1, \ldots, Q_L)^T$. First, it determines whether there are two distinct distributions Q_i, Q_j such that Q_i is not irreducible wrt Q_j (see Section 3 for definition of irreducibility), in which case $(Q_1, \ldots, Q_L)^T$ cannot be a permutation of $(P_1, \ldots, P_L)^T$. If there is such a pair, it returns a value indicating that $(Q_1, \ldots, Q_L)^T$ is not a permutation of $(W_1, \ldots, W_L)^T$. Otherwise, it finds for each Q_i the set of $\tilde{P}_j$ such that $\tilde{P}_j$ and Q_i lie on the same face of the simplex. Finally, using $\boldsymbol{S}$, it iteratively determines for each Q_i whether there is some P_j such that $P_j = Q_i$. If $(Q_1, \ldots, Q_L)^T$ is a permutation of $(P_1, \ldots, P_L)^T$, this procedure finds the appropriate permutation matrix; otherwise, the algorithm returns a value indicating that $(Q_1, \ldots, Q_L)^T$ is not a permutation of $(W_1, \ldots, W_L)^T$.

6 Estimation

In this section, we develop novel estimators that can be used to extend the Demix algorithm to the finite sample case. This will be the basis of finite sample algorithms for the demixing problem in mixed membership models and the decontamination problem of partial label models. Let $\mathcal{X} = \mathbb{R}^d$ be equipped with the standard Borel σ-algebra $\mathcal{C}$ and $\tilde{P}_1, \ldots, \tilde{P}_L$ be probability distributions on this space. Suppose that we observe for $i = 1, \ldots, L$,

$$X_1^i, \ldots, X_{n_i}^i \sim \tilde{P}_i.$$

Let $\mathcal{S}$ be any VC class with VC-dimension $V < \infty$, containing the set of all open balls, all open rectangles, or some other collection of sets that generates the Borel σ-algebra $\mathcal{C}$. Define $\epsilon_i(\delta_i) \equiv 3\sqrt{\frac{V \log(n_i+1) - \log(\delta_i/2)}{n_i}}$ for $i = 1, \ldots, L$.

Our goal is to establish estimators $\widehat{Q}_1, \ldots, \widehat{Q}_L$ that, when suitably permuted, converge uniformly on $\mathcal{S}$ to $P_1, \ldots, P_L$. Previously developed uniform convergence results assume access to i.i.d. samples and are based on the VC inequality (Blanchard et al., 2010). This inequality says that for each $i \in [L]$, and $\delta > 0$, the following holds with probability at least $1 - \delta$:

$$\sup_{S \in \mathcal{S}} |\tilde{P}_i(S) - \tilde{P}_i^\dagger(S)| \leqslant \epsilon_i(\delta)$$

where the *empirical distribution* is given by $\tilde{P}_i^\dagger(S) = \frac{1}{n_i} \sum_{j=1}^{n_i} \mathbf{1}_{\{X_j^i \in S\}}$. The challenge is that because of the recursive nature of the Demix algorithm, we cannot assume access to i.i.d. samples to estimate every distribution that arises. We show that uniform convergence of distributions propagates through the algorithm if we employ an estimator of κ^* with a known rate of convergence.

Let $\widehat{F}$ be an estimate of a distribution F and $\mathcal{D}(\widehat{F}) = \{i : \widehat{F}$ relies on data from the distribution $\tilde{P}_i\}$. We introduce the following estimator using the estimates $\widehat{F}$ and $\widehat{H}$:

$$\widehat{\kappa}(\widehat{F} \,|\, \widehat{H}) = \inf_{S \in \mathcal{S}} \frac{\widehat{F}(S) + \gamma(\mathcal{D}(\widehat{F}))}{(\widehat{H}(S) - \gamma(\mathcal{D}(\widehat{H})))_+}$$

where $\gamma(I) = \sum_{i \in I} \epsilon_i(\frac{1}{n_i})$ and $I \subseteq [L]$. Notice that when $\widehat{F}$ and $\widehat{H}$ are empirical distributions, e.g., $\widehat{F} = \tilde{P}_i^\dagger$ and $\widehat{H} = \tilde{P}_j^\dagger$, we recover the consistent estimator from Blanchard et al. (2010): $\widehat{\kappa}(\widehat{F} \,|\, \widehat{H}) = \inf_{S \in \mathcal{S}} \frac{\widehat{F}(S) + \epsilon_i(\frac{1}{n_i})}{(\widehat{H}(S) - \epsilon_j(\frac{1}{n_j}))_+}$.

Based on the estimator $\widehat{\kappa}$, we introduce the following estimator of the residue of F wrt H.

Algorithm 3 ResidueHat($\widehat{F} \,|\, \widehat{H}$)

Input: $\widehat{F}, \widehat{H}$ are estimates of F, H

1: $\widehat{\kappa} \longleftarrow \widehat{\kappa}(\widehat{F} \,|\, \widehat{H})$

2: **return** $\frac{\widehat{F} - \widehat{\kappa}(1 - \widehat{H})}{1 - \widehat{\kappa}}$

Definition 5. *Let $G \longleftarrow$ Residue$(F \,|\, H)$ and $\widehat{G} \longleftarrow$ ResidueHat$(\widehat{F} \,|\, \widehat{H})$. We call $\widehat{G}$ a* ResidueHat estimator *of G if (i) $F \neq H$, (ii) $F, H \in$ conv$(P_1, \ldots, P_L)$, and (iii) $\widehat{F}$ and $\widehat{H}$ are either empirical distributions or ResidueHat estimators of F and H.*

Note that the above definition is recursive and matches the recursive structure of the Demix algorithm.

To use ResidueHat estimators to estimate the P_is, we build on the rate of convergence result from Scott (2015). In Scott (2015), a rate of convergence was established for an estimator of κ^* using empirical distributions; we extend these results to our setting of recursive estimators and achieve the same rate of convergence. To ensure that this rate of convergence holds for every estimate in our algorithm, we introduce the following condition.

(A″) $P_1, \ldots, P_L$ are such that $\forall i \ \mathrm{supp}(P_i) \not\subseteq \cup_{j \neq i} \mathrm{supp}(P_j)$.

Note that this assumption is a natural generalization of the separability assumption.

Let $\boldsymbol{n} \equiv (n_1, \ldots, n_L)$; we write $\boldsymbol{n} \longrightarrow \infty$ to indicate that $\min_i n_i \longrightarrow \infty$. The following result establishes sufficient conditions under which ResidueHat estimates converge uniformly.

Proposition 5. *If $P_1, \ldots, P_L$ satisfy **(A″)** and $\widehat{G}$ is a ResidueHat estimator of a distribution $G \in \mathrm{conv}(P_1, \ldots, P_L)$, then $\sup_{S \in \mathcal{S}} |\widehat{G}(S) - G(S)| \xrightarrow{i.p.} 0$ as $\boldsymbol{n} \longrightarrow \infty$.*

Sketch. Let $\widehat{F}_1, \ldots, \widehat{F}_K$ denote the ResidueHat estimators in terms of which $\widehat{G}$ is defined and $F_1, \ldots, F_K$ the distributions which they estimate. We show that if $\widehat{F}_i$ and $\widehat{F}_j$ satisfy uniform deviation inequalities, $F_i \neq F_j$, and $\widehat{F}_l \longleftarrow \mathrm{ResidueHat}(\widehat{F}_i \,|\, \widehat{F}_j)$, then there exists constants $A_1, A_2 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{F}_i) \cup \mathcal{D}(\widehat{F}_j))$, with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{F}_i) \cup \mathcal{D}(\widehat{F}_j)} \frac{1}{n_i}$ and for all $S \in \mathcal{S}$,

$$|\widehat{\kappa}(\widehat{F}_i \,|\, \widehat{F}_j) - \kappa^*(F_i \,|\, F_j)| \leqslant A_1 \gamma(\mathcal{D}(\widehat{F}_i) \cup \mathcal{D}(\widehat{F}_j))$$
$$|\widehat{F}_l(S) - F_l(S)| < A_1 \gamma(\mathcal{D}(\widehat{F}_l))$$

Since $F_i \neq F_j$, $\widehat{F}_i$ and $\widehat{F}_j$ satisfy uniform deviation inequalites by assumption, and $P_1, \ldots, P_L$ satisfy **(A″)**, we show that by Lemma 12, the first inequality (a rate of convergence) holds. Then, by the first inequality, $F_i \neq F_j$, and the assumption that $\widehat{F}_i$ and $\widehat{F}_j$ satisfy uniform deviation inequalites, we have that Lemma 10 implies the second inequality. Then, reasoning inductively, we obtain the result. □

Based on the ResidueHat estimators, we introduce an empirical version of the Demix algorithm, namely, DemixHat (see Algorithm 11). The only substantial difference is that we replace the Residue function with the ResidueHat function. See the Appendix for details.

We now state our main estimation result.

Theorem 4. *Let $\epsilon > 0$. Suppose that $P_1, \ldots, P_L$ satisfy **(A″)** and $\boldsymbol{\Pi}$ satisfies **(B′)**. Then, with probability tending to 1 as $\boldsymbol{n} \longrightarrow \infty$, DemixHat($\widehat{P}_1^\dagger, \ldots, \widehat{P}_L^\dagger$) returns $(\widehat{Q}_1, \ldots, \widehat{Q}_L)$ and there exists a permutation $\sigma : [L] \longrightarrow [L]$ such that for every $i \in [L]$,*

$$\sup_{S \in \mathcal{S}} |\widehat{Q}_i(S) - P_{\sigma(i)}(S)| < \epsilon.$$

In the preceding, we have assumed a fixed VC class to simplify the presentation. However, these results easily extend to the setting where $\mathcal{S} = \mathcal{S}_k$ and $k \longrightarrow \infty$ at a suitable rate depending on the growth of the VC dimensions V_k. This allows for the P_is to be estimated uniformly on arbitrarily complex events, e.g., $\mathcal{S}_k$ is the set of unions of k open balls.

7 Discussion

In the Appendix, we present a finite sample algorithm for the decontamination of a partial label model (see Algorithm 5). This algorithm is based on a different approach from the PartialLabel Algorithm 2: it combines DemixHat with an empirical version of the VertexTest algorithm (see Algorithm 13). The reason for this hinges in the advantages and disadvantages associated with the two-sample κ^* operator and multi-sample κ^* operator, respectively. Algorithms that only use the two-sample κ^* operator have the following two advantages: (i) the geometry of the 2-sample κ^* operator is simpler than the geometry of the multi-sample κ^* operator and, as such, can be more tractable. Indeed, in recent years, several practical algorithms for estimating the two-sample $\kappa*$ have been developed (see Jain et al. (2016) and references therein). (ii) We have estimators with established rates of convergence for the two-sample κ^* operator, but not for the multi-sample κ^* operator. On the other hand, algorithms that use the multi-sample κ^* operator can be simpler and have the potential to be more practical since they can reduce the number of estimation steps.

8 Appendix

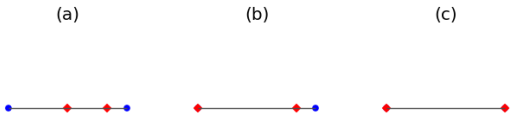

Figure 1: In (a), we are given a demixing problem where $L = 2$. The diamonds represent the mixture proportions of $\tilde{P}_1$ and $\tilde{P}_2$. The circles represent the base distributions. In (b), the residue of a contaminated distribution wrt the other contaminated distribution is computed (line 3), yielding a base distribution. In (c), the residue is computed again switching the roles of the contaminated distributions (line 4); this yields the remaining base distribution.

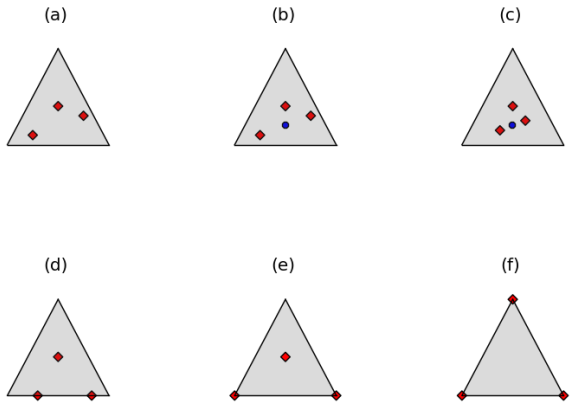

Figure 2: In (a), we are given a demixing problem with $M = L = 3$. The diamonds represent the mixture proportions of $\tilde{P}_1, \tilde{P}_2$ and $\tilde{P}_3$. In (b), the blue circle is a random distribution chosen in the convex hull of two of the distributions (line 7). In (c), two of the distributions are resampled so that their residues wrt the other distribution are on the same face of the simplex (lines 12-15). In (d), these particular residues are computed (lines 12-15). In (e), two of the distributions are demixed (lines 3-5). In (f), the residue of the final distribution wrt the final two demixed distributions is computed to obtain the final demixing (line 18-21).

8.1 Notation

Let A be a set. A° denotes the relative interior of A, i.e., $A^\circ = \{x \in A | B(x, r) \cap \text{aff}\, A \subseteq A \text{ for some } r > 0\}$. Then, ∂A denotes the relative boundary of A, i.e., $\partial A = A \backslash A^\circ$. In addition, let $\|\cdot\|$ denote an arbitrary norm.

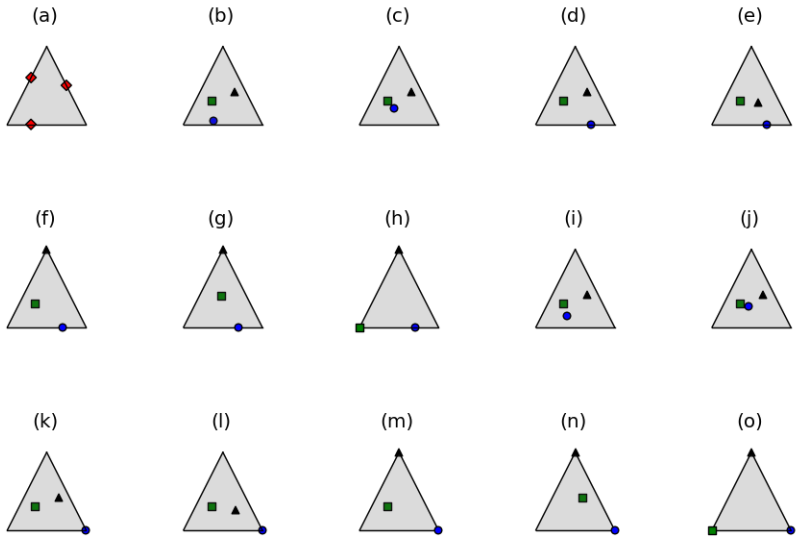

Figure 3: We are given an instance of a partial label problem where $M = L = 3$ and the partial labels each have two 1s and a single 0. In (a), the red diamonds represent the mixture proportions of the distributions $\tilde{P}_1, \tilde{P}_2, \tilde{P}_3$. In (b), three distributions Q_1, Q_2, Q_3 are sampled uniformly randomly from the convex hull of $\tilde{P}_1, \tilde{P}_2, \tilde{P}_3$; the green square, black triangle, and blue circle represent their mixture proportions. Figures (c)-(h) show how the algorithm generates a set of candidate distributions $(W_1^{(2)}, W_2^{(2)}, W_3^{(2)})^T$ with $k = 2$. In (h), PartialLabel runs VertexTest on $(W_1^{(2)}, W_2^{(2)}, W_3^{(2)})^T$ and determines that $(W_1^{(2)}, W_2^{(2)}, W_3^{(2)})^T$ is not a permutation of $(P_1, P_2, P_3)^T$. In (i)-(o), PartialLabel begins again with Q_1, Q_2, Q_3 and executes the same series of steps with $k = 3$, generating $(W_1^{(3)}, W_2^{(3)}, W_3^{(3)})^T$. In (o), it runs VertexTest on $(W_1^{(3)}, W_2^{(3)}, W_3^{(3)})^T$ and determines that $(W_1^{(3)}, W_2^{(3)}, W_3^{(3)})^T$ is a permutation of $(P_1, P_2, P_3)^T$.

For two vectors, $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^K$, define

$$\min(\boldsymbol{x}^T, \boldsymbol{y}^T) = (\min(x_1, y_1), \ldots, \min(x_K, y_K))$$

We use the following affine mapping throughout the paper: $m_\nu(\boldsymbol{x}, \boldsymbol{y}) = (1 - \nu)\boldsymbol{x} + \nu\boldsymbol{y}$ where $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^L$ and $\nu \in [0, 1]$. Overloading notation, when Q_1 and Q_2 are distributions, we define $m_\nu(Q_1, Q_2) = (1 - \nu)Q_1 + \nu Q_2$. We call ν the *resampling proportion*. Note that if $\boldsymbol{\eta}_1, \boldsymbol{\eta}_2 \in \Delta_L$ and $Q_1 = \boldsymbol{\eta}_1^T \boldsymbol{P}$ and $Q_2 = \boldsymbol{\eta}_2^T \boldsymbol{P}$, then $m_\nu(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2)$ is the mixture proportion for the distribution $m_\nu(Q_1, Q_2)$.

8.2 Necessary Conditions

8.2.1 Mutual Contamination Models

Theorem 1. **(A)** We prove the contrapositive. Let ϕ be a demixing operator and $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. Suppose that there is some P_i and $Q \in \text{conv}(\{P_j : j \neq i\})$ with $Q = \sum_{j \neq i} \beta_j P_j$ such that P_i is not irreducible wrt Q. Then, there is some distribution G and $\gamma \in (0, 1]$ such that $P_i = \gamma Q + (1 - \gamma)G$.

Suppose $\gamma = 1$. Then, $P_i = Q \in \text{conv}(\{P_k : k \neq i\})$. But, then $\tilde{P}_1, \ldots, \tilde{P}_M \in \text{conv}(\{P_j : j \neq i\} \cup \{R\})$ for any distribution $R \notin \text{conv}(P_1, \ldots, P_L)$. This shows that ϕ does not satisfy **(M)**.

Therefore, assume that $\gamma \in (0, 1)$. Either $G \in \text{conv}(P_1, \ldots, P_L)$ or $G \notin \text{conv}(P_1, \ldots, P_L)$. Suppose that $G \in \text{conv}(P_1, \ldots, P_L)$. Then, there exist $\alpha_1, \ldots, \alpha_L$ all nonnegative and summing to 1 such that

$$P_i = \gamma Q + (1 - \gamma)(\alpha_1 P_1 + \ldots + \alpha_L P_L)$$

Therefore, $P_i \in \text{conv}(\{P_k : k \neq i\})$. But then $\tilde{P}_1, \ldots, \tilde{P}_M \in \text{conv}(\{P_j : j \neq i\} \cup \{R\})$ for any distribution $R \notin \text{conv}(P_1, \ldots, P_L)$. This shows that that ϕ does not satisfy **(M)**.

Now, suppose that $G \notin \text{conv}(P_1, \ldots, P_L)$. Since $P_i \in \text{conv}(G, Q)$ and $Q \in \text{conv}(\{P_j : j \neq i\})$, we have that $\text{conv}(\{P_j : j \neq i\} \cup \{G\}) \supset \text{conv}(P_1, \ldots, P_L)$. Then, $\tilde{P}_1, \ldots, \tilde{P}_M \in \text{conv}(\{P_j : j \neq i\} \cup \{G\})$. This shows that ϕ does not satisfy **(M)**. The result follows.

(B) Let $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. Clearly, $\dim \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) \leqslant \dim \text{span}(P_1, \ldots, P_L)$ since $\tilde{P}_i = \boldsymbol{\pi}_i^T \boldsymbol{P}$ for all $i \in [M]$. Since ϕ satisfies **(L)**, $\text{span}(P_1, \ldots, P_L) \subset \text{span}(\tilde{P}_1, \ldots, \tilde{P}_L)$, which implies that $\dim \text{span}(P_1, \ldots, P_L) \leqslant \dim \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M)$. Therefore, $\dim \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) = \dim \text{span}(P_1, \ldots, P_L)$. Then, since $\tilde{P}_1, \ldots, \tilde{P}_M \in \text{range}(\boldsymbol{\Pi})$, $\dim \text{span}(P_1, \ldots, P_L) \leqslant \dim \text{range}\, \boldsymbol{\Pi}$. By Result 3.117 of Axler (2015), $\text{rank}(\boldsymbol{\Pi}) = \dim \text{range}\, \boldsymbol{\Pi} \geqslant \dim \text{span}(P_1, \ldots, P_L)$.

□

Proposition 2. Suppose that $\phi(\tilde{\boldsymbol{P}}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. We first show that ϕ satisfies **(L)**. By hypothesis, $P_1, \ldots, P_L$ are jointly irreducible. By Lemma B.1 of Blanchard and Scott (2014), $P_1, \ldots, P_L$ are linearly independent. Since by hypothesis $\boldsymbol{\Pi}$ has full rank, there exist L rows in $\boldsymbol{\Pi}$, $\boldsymbol{\pi}_{i_1}, \ldots, \boldsymbol{\pi}_{i_L}$, that are linearly independent. By Lemma B.1 of Blanchard and Scott (2014), $\tilde{P}_{i_1}, \ldots, \tilde{P}_{i_L}$ are linearly independent. Since $\boldsymbol{\Pi}\boldsymbol{P} = \tilde{\boldsymbol{P}}$, $\text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) \subseteq \text{span}(P_1, \ldots, P_L)$. But, then since $\dim \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) \geqslant L$, we have $\text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) = \text{span}(P_1, \ldots, P_L)$. Therefore, ϕ satisfies **(L)**.

Now, we show that ϕ satisfies **(M)**. Suppose that there is another solution $(\boldsymbol{\Pi}', \boldsymbol{P}')$ with $\boldsymbol{P}' = (P_1, \ldots, P_L)^T$ such that $\boldsymbol{\Pi}'\boldsymbol{P}' = \tilde{\boldsymbol{P}}$ and with some P_i' such that $P_i' \notin \text{conv}(P_1, \ldots, P_L)$. We claim that $P_i' \notin \text{span}(P_1, \ldots, P_L)$. If $P_i' \in \text{span}(P_1, \ldots, P_L)$, then we must have that $P_i' = \sum_{i=1}^L a_i P_i$ where at least one of the a_i is negative. But, by joint irreducibility of $P_1, \ldots, P_L$, P_i' is not a distribution, which is a contradiction. So, the claim follows. But, then, since $\text{span}(P_1, \ldots, P_L) = \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M)$, we must have that $\text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) \subseteq \text{span}(P_1', \ldots, P_{i-1}', P_{i+1}', \ldots, P_L)$, which is impossible since $\dim \text{span}(\tilde{P}_1, \ldots, \tilde{P}_M) = L$. □

8.2.2 Partial Label Model

Definition 6. *A decontamination operator ψ satisfies Maximality **(M′)** iff $\psi(\tilde{\boldsymbol{P}}, \boldsymbol{S}) = (\boldsymbol{\Pi}, \boldsymbol{P})$ implies that any pair $(\boldsymbol{\Pi}', \boldsymbol{P}')$ with $\boldsymbol{\Pi}' \in \Delta_L^M$ and $\boldsymbol{P}' = (P_1', \ldots, P_L')^T \in \mathcal{P}^L$ that solves (1) and is consistent with $\boldsymbol{S}$ is such that $\{P_1', \ldots, P_L'\} \subseteq \text{conv}(P_1, \ldots, P_L)$.*

The notion of linearity for a decontamination operator for the partial label model is identical to the notion of linearity for a decontamination operator for a mixed membership model.

Definition 7. *A decontamination operator ψ satisfies Linearity **(L′)** iff $\psi_2(\tilde{\boldsymbol{P}}, \boldsymbol{S}) = \boldsymbol{P}$ implies $\{P_1, \ldots, P_L\} \subseteq$* span$(\tilde{P}_1, \ldots, \tilde{P}_M)$.

Proposition 6. *Let ψ denote a decontamination operator for the partial label model and $\psi(\tilde{\boldsymbol{P}}, \boldsymbol{S}) = (\boldsymbol{\Pi}, \boldsymbol{P})$. (i) If ψ satisfies **(L′)**, then $(\boldsymbol{\Pi}, \boldsymbol{P})$ need not satisfy **(A)**. (ii) If ψ satisfies **(L′)**, then $(\boldsymbol{\Pi}, \boldsymbol{P})$ satisfies **(B)**.*

Proof. **(i)** Let $Q_1 \sim$ unif$(0, 2)$ and $Q_2 \sim$ unif$(1, 3)$. Let $P_1 = \frac{2}{3}Q_1 + \frac{1}{3}Q_2$, $P_2 = \frac{1}{3}Q_1 + \frac{2}{3}Q_2$, $\tilde{P}_1 = P_1$, and $\tilde{P}_2 = P_2$. Then, $\boldsymbol{S} = \boldsymbol{I}_2$, identity matrix. Then, any $(\boldsymbol{\Pi}', \boldsymbol{P}')$ that satisfies (1) and is consistent with $\boldsymbol{S}$ must be such that $(P_1, P_2)^T = \boldsymbol{P}'$. Therefore, **(M′)** is satisfied. But, clearly, **(A)** is not satisfied.

(ii) The proof is identical to the proof that **(L)** implies **(B)** for mixed membership models.

□

Proposition 1. We give a proof by contraposition. Suppose that there exists $i \neq j$ such that $\boldsymbol{S}_{:,i} = \boldsymbol{S}_{:,j}$. Without loss of generality, let $i = 1$ and $j = 2$. Suppose that $(\boldsymbol{\Pi}, \boldsymbol{P})$ is consistent with $\boldsymbol{S}$ and solves $\tilde{\boldsymbol{P}} = \boldsymbol{\Pi}\boldsymbol{P}$. Then, the pair $(\boldsymbol{\Pi}', \boldsymbol{P}')$

$$\boldsymbol{\Pi}' = \begin{pmatrix} \boldsymbol{\Pi}_{:,2} & \boldsymbol{\Pi}_{:,1} & \boldsymbol{\Pi}_{:,3} & \ldots & \boldsymbol{\Pi}_{:,L} \end{pmatrix}$$

$$\boldsymbol{P}' = \begin{pmatrix} P_2 \\ P_1 \\ P_3 \\ \vdots \\ P_L \end{pmatrix}$$

solves $\tilde{\boldsymbol{P}} = \boldsymbol{\Pi}'\boldsymbol{P}'$ and is consistent with $\boldsymbol{S}$. □

Proposition 7. *If a decontamination operator ψ is such that $\psi((\tilde{\boldsymbol{P}}, \boldsymbol{S})) = (\boldsymbol{\Pi}, \boldsymbol{P})$ only if $(\boldsymbol{\Pi}, \boldsymbol{P})$ satisfy **(A′)** and **(B′)**, then ϕ satisfies **(L)**.*

Proof. The proof is identical to the first paragraph of the proof of Proposition 2. □

8.3 Residue Algorithms

Algorithm 4 Residue$(F_0 \mid F_1)$

1: $\kappa \longleftarrow \kappa^*(F_0 \mid F_1)$
2: **return** $\frac{F_0 - \kappa(1-F_1)}{1-\kappa}$

Algorithm 5 Residue$(F_0 \mid \{F_1, \ldots, F_K\})$

1: $(\nu_1, \ldots, \nu_K) \longleftarrow (\nu'_1, \ldots, \nu'_K)$ achieving the maximum in $\kappa^*(F_0 \mid F_1, \ldots, F_K)$
2: **return** $\frac{F_0 - \sum_{i=1}^K \nu_i F_i}{1 - \sum_{i=1}^K \nu_i}$

8.4 General Lemmas

Prop 4. Without loss of generality, suppose $l = 1$ and let $A = [L]\backslash\{1\}$; the other cases follow by a similar argument. Suppose G is such that

$$Q_1 = (1 - \sum_{j \geqslant 2} \mu_j)G + \sum_{j \geqslant 2} \mu_j Q_j$$

Note that since $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_L$ are linearly independent and $P_1, \ldots, P_L$ are jointly irreducible, $Q_1, \ldots, Q_L$ are linearly independent by Lemma B.1 from Blanchard and Scott (2014). Therefore, $\sum_{j \geqslant 2} \mu_j < 1$ because, if not, $Q_1 = \sum_{j \geqslant 2} \mu_j Q_j$. Therefore, $\kappa^*(Q_1 \,|\, \{Q_j : j \neq 1\}) < 1$.

Further, any G that satisfies the above equation has the form $\sum_{i=1}^{L} \gamma_i P_i$. The γ_i must sum to one, and we have that they are nonnegative by joint irreducibility. That is, $\boldsymbol{\gamma} \equiv \left(\gamma_1, \ldots, \gamma_L\right)^T$ is a discrete distribution. Then, the above equation is equivalent to

$$\boldsymbol{\eta}_1^T \boldsymbol{P} = (1 - \sum_{j \geqslant 2} \mu_j)\boldsymbol{\gamma}^T \boldsymbol{P} + \sum_{j \geqslant 2} \mu_j \boldsymbol{\eta}_j^T \boldsymbol{P}$$

Since $P_1, \ldots, P_L$ are jointly irreducible, $P_1, \ldots, P_L$ are linearly independent by Lemma B.1 Blanchard and Scott (2014). By linear independence of $P_1, \ldots, P_L$, we obtain

$$\boldsymbol{\eta}_1 = (1 - \sum_{j \geqslant 2} \mu_j)\boldsymbol{\gamma} + \sum_{j \geqslant 2} \mu_j \boldsymbol{\eta}_j$$

Consequently, $\kappa^*(Q_1 \,|\, \{Q_j : j \neq 1\}) = \kappa^*(\boldsymbol{\eta}_1 \,|\, \{\boldsymbol{\eta}_j : j \neq 1\}) < 1$ and there is a one-to-one correspondence between the feasible solution in $\kappa^*(Q_1 \,|\, \{Q_j : j \neq 1\})$ and the feasible solution in $\kappa^*(\boldsymbol{\eta}_1 \,|\, \{\boldsymbol{\eta}_j : j \neq 1\})$. The one-to-one correspondence is given by $G = \boldsymbol{\gamma}^T \boldsymbol{P}$.

There is at least one point attaining the maximum in the optimization problem $\kappa^*(\boldsymbol{\eta}_1 \,|\, \{\boldsymbol{\eta}_j : j \neq 1\})$ by Lemma A.1 in Blanchard and Scott (2014).

We see that the maximizing μ_j are unique as follows. Suppose

$$Q_1 = (1 - \kappa^*)G + \sum_{j \geqslant 2} \mu_j Q_j = (1 - \kappa^*)G + \sum_{j \geqslant 2} \mu'_j Q_j$$

The linear independence of $Q_1, \ldots, Q_L$ implies that $\mu_j = \mu'_j$. □

Lemma 1 gives us some useful properties of the two-sample κ^* that we exploit in the PartialLabel and Demix algorithms. Statement *1* gives an alternative form of κ^*. Statement *2* gives the intuitive result that the residues lie on the boundary of the simplex. Statement *3* gives a useful relation for determining whether two mixture proportions are on the same face; we use this relation extensively in our algorithms.

Lemma 1. *Let $F_1, \ldots, F_K$ be jointly irreducible distributions with $\boldsymbol{F} = (F_1, \ldots, F_K)^T$, Q_1, Q_2 be two distributions such that $Q_i = \boldsymbol{\eta}_i^T \boldsymbol{F}$ where $\boldsymbol{\eta}_i \in \Delta_K$ for $i = 1, 2$ and $\boldsymbol{\eta}_1 \neq \boldsymbol{\eta}_2$. Let R be the residue of Q_1 wrt Q_2 and $R = \boldsymbol{\mu}^T \boldsymbol{F}$.*

1. *We have the following equivalence of optimization problems:*

$$\kappa^*(Q_1 \,|\, Q_2) = \max(\alpha \geqslant 1 | \exists \text{ a distribution } G \text{ s.t } G = Q_2 + \alpha(Q_1 - Q_2))$$

2. $\boldsymbol{\mu} \in \partial\Delta_K$.

3. $\mathcal{S}(\boldsymbol{\eta}_2) \nsubseteq \mathcal{S}(\boldsymbol{\eta}_1)$ *if and only if* $R = Q_1$ *if and only if* $\kappa^*(Q_1 \,|\, Q_2) = 0$.

Lemma 1. 1. Consider the linear relation: $Q_1 = (1-\kappa)G + \kappa Q_2$ where $\kappa \in [0, 1]$. Since $F_1, \ldots, F_K$ are jointly irreducible and $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_2$ are linearly independent, Q_1 and Q_2 are linearly independent by Lemma B.1 of Blanchard and Scott (2014). Therefore, $\kappa < 1$. We can rewrite the relation as follows:

$$G = \frac{1}{1-\kappa} Q_1 - \frac{\kappa}{1-\kappa} Q_2 = \alpha Q_1 + (1-\alpha) Q_2$$

where $\alpha = \frac{1}{1-\kappa}$. The equivalence follows.

2. Since R is the residue of Q_1 wrt Q_2, by Proposition 4, $\boldsymbol{\mu}$ is the residue of $\boldsymbol{\eta}_1$ wrt $\boldsymbol{\eta}_2$ and $\boldsymbol{\mu} \in \Delta_K$. Therefore, by statement *1* in Lemma 1, $\boldsymbol{\mu}$ is such that α^* is maximized subject to the following constraints:

$$\begin{aligned} \boldsymbol{\mu} &= (1-\alpha^*)\boldsymbol{\eta}_2 + \alpha^* \boldsymbol{\eta}_1 \\ \alpha^* &\geqslant 1 \\ \boldsymbol{\mu} &\in \Delta_K \end{aligned}$$

Suppose that $\min_i \mu_i > 0$. Then, there is some $\epsilon > 0$ such that

$$\begin{aligned} \boldsymbol{\mu}' &= (1-\alpha^*-\epsilon)\boldsymbol{\eta}_2 + (\alpha^*+\epsilon)\boldsymbol{\eta}_1 \\ \alpha^* + \epsilon &\geqslant 1 \\ \boldsymbol{\mu}' &\in \Delta_K \end{aligned}$$

But, this contradicts the definition of α^* and $\boldsymbol{\mu}$. Therefore, $\min_i \mu_i = 0$. Consequently, $\boldsymbol{\mu} \in \partial\Delta_K$.

3. By definition of κ^*, it is clear that $R = Q_1$ if and only if $\kappa^*(Q_1 \,|\, Q_2) = 0$. Therefore, it suffices to show that $\mathcal{S}(\boldsymbol{\eta}_2) \nsubseteq \mathcal{S}(\boldsymbol{\eta}_1)$ if and only if $\kappa^*(Q_1 \,|\, Q_2) = 0$. Suppose $\mathcal{S}(\boldsymbol{\eta}_2) \nsubseteq \mathcal{S}(\boldsymbol{\eta}_1)$. Then, there must be $i \in [K]$ such that $\eta_{2,i} > 0$ and $\eta_{1,i} = 0$. Then, for $\alpha > 1$,

$$\min_{i \in [K]} (1-\alpha)\eta_{2,i} + \alpha \eta_{1,i} < 0$$

But, this violates the constraint of the optimization problem. Therefore, $\alpha = 1$. By statement *1* in Lemma 1, $\kappa^*(\boldsymbol{\eta}_1 \,|\, \boldsymbol{\eta}_2) = 0$. By Proposition 4, $\kappa^*(Q_1 \,|\, Q_2) = \kappa^*(\boldsymbol{\eta}_1 \,|\, \boldsymbol{\eta}_2) = 0$.

Now, suppose $\mathcal{S}(\boldsymbol{\eta}_2) \subseteq \mathcal{S}(\boldsymbol{\eta}_1)$. Then, for any $i \in [K]$, if $\eta_{2,i} > 0$, then $\eta_{1,i} > 0$. Then, there is $\alpha > 1$ sufficiently close to 1 such that

$$\min_{i \in [K]} \eta_{2,i} + \alpha(\eta_{1,i} - \eta_{2,i}) \geqslant 0$$

By statement *1* in this Lemma, $\kappa^*(\boldsymbol{\eta}_1 \,|\, \boldsymbol{\eta}_2) > 0$. By proposition 4, $\kappa^*(Q_1 \,|\, Q_2) = \kappa^*(\boldsymbol{\eta}_1 \,|\, \boldsymbol{\eta}_2) > 0$. □

8.5 Demixing Mixed Membership Models

8.5.1 Lemmas

Lemma 2 establishes an intuitive continuity property of the two-sample version of κ^* and the residue.

Lemma 2. *Let $\boldsymbol{\eta}_1, \boldsymbol{\eta}_2, \boldsymbol{\eta}_3 \in \Delta_L$ be distinct vectors and let $\boldsymbol{\mu}$ be the residue of $\boldsymbol{\eta}_2$ wrt $\boldsymbol{\eta}_1$. Let α_n be a sequence such that $\alpha_n \longrightarrow 1$ as $n \longrightarrow \infty$ and let $\boldsymbol{\tau}_n$ be the residue of $\alpha_n\boldsymbol{\eta}_2 + (1-\alpha_n)\boldsymbol{\eta}_3$ wrt $\boldsymbol{\eta}_1$. Then,*

1. $\lim_{n\longrightarrow\infty} \kappa^*(\alpha_n\boldsymbol{\eta}_2 + (1-\alpha_n)\boldsymbol{\eta}_3 \,|\, \boldsymbol{\eta}_1) = \kappa^*(\boldsymbol{\eta}_2 \,|\, \boldsymbol{\eta}_1)$
2. $\lim_{n\longrightarrow\infty} \|\boldsymbol{\tau}_n - \boldsymbol{\mu}\| = 0$

Lemma 2. 1. In order to apply the residue operator κ^* to $\boldsymbol{\eta}_1, \boldsymbol{\eta}_2, \boldsymbol{\eta}_3$, we think of $\boldsymbol{\eta}_1, \boldsymbol{\eta}_2, \boldsymbol{\eta}_3$ as discrete probability distributions. Let $(\mathcal{X}, \mathcal{C})$ be the measurable space on which $\boldsymbol{\eta}_1, \boldsymbol{\eta}_2, \boldsymbol{\eta}_3$ are defined. Note that $\mathcal{X}$ is finite and let $\mathcal{C} = 2^{\mathcal{X}}$ (i.e., the power set of $\mathcal{X}$), which is also finite. By Proposition 1 of Blanchard and Scott (2014),

$$\kappa^*(\boldsymbol{\eta}_2 \,|\, \boldsymbol{\eta}_1) = \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\boldsymbol{\eta}_2(C)}{\boldsymbol{\eta}_1(C)}$$

Since $\mathcal{C}$ is finite, there is $\epsilon > 0$ such that $\inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \boldsymbol{\eta}_1(C) > \epsilon$. Then,

$$\begin{aligned}
\kappa^*(\alpha_n\boldsymbol{\eta}_2 + (1-\alpha_n)\boldsymbol{\eta}_3 \,|\, \boldsymbol{\eta}_1) &= \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\alpha_n\boldsymbol{\eta}_2(C) + (1-\alpha_n)\boldsymbol{\eta}_3(C)}{\boldsymbol{\eta}_1(C)} \\
&\leqslant \alpha_n \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\boldsymbol{\eta}_2(C)}{\boldsymbol{\eta}_1(C)} + \frac{1-\alpha_n}{\epsilon} \\
&\longrightarrow \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\boldsymbol{\eta}_2(C)}{\boldsymbol{\eta}_1(C)}
\end{aligned}$$

as $n \longrightarrow \infty$. Further, since $\frac{\boldsymbol{\eta}_2(\cdot)}{\boldsymbol{\eta}_1(\cdot)}$ and $\frac{\boldsymbol{\eta}_3(\cdot)}{\boldsymbol{\eta}_1(\cdot)}$ are bounded over $\{C \in \mathcal{C} : \boldsymbol{\eta}_1(C) > 0\}$,

$$\begin{aligned}
\kappa^*(\alpha_n\boldsymbol{\eta}_2 + (1-\alpha_n)\boldsymbol{\eta}_3 \,|\, \boldsymbol{\eta}_1) &= \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\alpha_n\boldsymbol{\eta}_2(C) + (1-\alpha_n)\boldsymbol{\eta}_3(C)}{\boldsymbol{\eta}_1(C)} \\
&\geqslant \alpha_n \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\boldsymbol{\eta}_2(C)}{\boldsymbol{\eta}_1(C)} + (1-\alpha_n) \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\boldsymbol{\eta}_3(C)}{\boldsymbol{\eta}_1(C)} \\
&\longrightarrow \inf_{C\in\mathcal{C}, \boldsymbol{\eta}_1(C)>0} \frac{\boldsymbol{\eta}_2(C)}{\boldsymbol{\eta}_1(C)}
\end{aligned}$$

as $n \longrightarrow \infty$. By the sandwich principle of limits, the result follows.

2. We write $\boldsymbol{\mu} = \kappa\boldsymbol{\eta}_2 + (1-\kappa)\boldsymbol{\eta}_1$ and $\boldsymbol{\tau}_n = \kappa_n(\alpha_n\boldsymbol{\eta}_2 + (1-\alpha_n)\boldsymbol{\eta}_3) + (1-\kappa_n)\boldsymbol{\eta}_1$ where $\kappa = \kappa^*(\boldsymbol{\eta}_2 \,|\, \boldsymbol{\eta}_1)$ and $\kappa_n = \kappa^*(\alpha_n\boldsymbol{\eta}_2 + (1-\alpha_n)\boldsymbol{\eta}_3 \,|\, \boldsymbol{\eta}_1)$. Then, by the triangle inequality,

$$\begin{aligned}
\|\boldsymbol{\mu} - \boldsymbol{\tau}_n\| &= \|(\kappa_n - \kappa)\boldsymbol{\eta}_1 + (\kappa - \kappa_n\alpha_n)\boldsymbol{\eta}_2 - (1-\alpha_n)\boldsymbol{\eta}_3\| \\
&\leqslant |\kappa_n - \kappa| \, \|\boldsymbol{\eta}_1\| + |\kappa - \kappa_n\alpha_n| \, \|\boldsymbol{\eta}_2\| + |1-\alpha_n| \, \|\boldsymbol{\eta}_3\| \\
&\longrightarrow 0
\end{aligned}$$

as $n \longrightarrow \infty$ since $\alpha_n \longrightarrow 1$ and $\kappa_n \longrightarrow \kappa$ by statement 1 in Lemma 2.

□

Lemma 3 guarantees that certain operations in the Demix algorithm preserve linear independence of the mixture proportions.

Lemma 3. *Let $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_K \in \Delta_L$ be linearly independent and $P_1, \ldots, P_L$ be jointly irreducible. Let $Q_i = \boldsymbol{\tau}_i^T \boldsymbol{P}$ for $i \in [K]$. Then for any $i, j \in [K]$ such that $i \neq j$,*

1. *If $\boldsymbol{\eta} = \sum_{k=1}^K a_k \boldsymbol{\tau}_k$ with $a_j \neq 0$, then $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_{j-1}, \boldsymbol{\eta}, \boldsymbol{\tau}_{j+1}, \ldots, \boldsymbol{\tau}_K$ are linearly independent.*
2. *Let R_k be the residue of Q_k with respect to Q_j for all $k \in [K] \backslash \{j\}$. Then, $R_k = \boldsymbol{\eta}_k^T \boldsymbol{P}$ where $\boldsymbol{\eta}_k \in \Delta_K$ and $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_{j-1}, \boldsymbol{\tau}_j, \boldsymbol{\eta}_{j+1}, \ldots, \boldsymbol{\eta}_K$ are linearly independent.*
3. *Let $\boldsymbol{\tau}^* \in \text{conv}(\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_k)^\circ$ and $\boldsymbol{\eta}_i \in \text{conv}(\boldsymbol{\tau}_i, \boldsymbol{\tau}^*)^\circ$ for $i \in [k]$ where $k \leqslant K$. Then,*

$$\boldsymbol{\eta}_1, \boldsymbol{\eta}_2, \ldots, \boldsymbol{\eta}_k, \boldsymbol{\tau}_{k+1}, \ldots, \boldsymbol{\tau}_K$$

are linearly independent.

Lemma 3.

1. Let $\boldsymbol{\eta} = \sum_{k=1}^K a_k \boldsymbol{\tau}_k$ with $a_j \neq 0$ for some $j \in [K]$. Let $b_1, \ldots, b_K \in \mathbb{R}$ such that

$$\begin{aligned} 0 &= b_1 \boldsymbol{\tau}_1 + \ldots + b_{j-1} \boldsymbol{\tau}_{j-1} + b_j \boldsymbol{\eta} + b_{j+1} \boldsymbol{\tau}_{j+1} + \ldots + b_K \boldsymbol{\tau}_K \\ &= (b_1 + b_j a_1) \boldsymbol{\tau}_1 + \ldots + (b_{j-1} + b_j a_{j-1}) \boldsymbol{\tau}_{j-1} + b_j a_j \boldsymbol{\tau}_j + \\ &(b_{j+1} + b_j a_{j+1}) \boldsymbol{\tau}_{j+1} + \ldots + (b_K + b_j a_K) \boldsymbol{\tau}_K \end{aligned}$$

   Since $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_K$ are linearly independent, $b_j a_j = 0$. Since $a_j \neq 0$ by hypothesis, $b_j = 0$. Therefore, the previous equation reduces to

$$b_1 \boldsymbol{\tau}_1 + \ldots + b_{j-1} \boldsymbol{\tau}_{j-1} + b_{j+1} \boldsymbol{\tau}_{j+1} + \ldots + b_K \boldsymbol{\tau}_K = 0$$

   By the linear independence of $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_K$, $b_1 = \ldots = b_K = 0$. The result follows.

2. Let $k \in [K] \backslash \{j\}$. By Proposition 4, $\boldsymbol{\eta}_k$ is the residue of $\boldsymbol{\tau}_k$ with respect to $\boldsymbol{\tau}_j$ and $\boldsymbol{\eta}_k \in \Delta_L$. By statement 1 of Lemma 1, there is $\alpha_k \geqslant 1$ such that $\boldsymbol{\eta}_k = (1 - \alpha_k) \boldsymbol{\tau}_j + \alpha \boldsymbol{\tau}_k$. By statement 1 of Lemma 3, $\boldsymbol{\eta}_1, \boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_k$ are linearly independent. Using induction and statement 1 of 3, $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_{j-1}, \boldsymbol{\tau}_j, \boldsymbol{\eta}_{j+1}, \ldots, \boldsymbol{\eta}_k$ are linearly independent.

3. We show the result inductively. By statement *1* of Lemma 3, it follows immediately that $\boldsymbol{\eta}_1, \boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_K$ are linearly independent. Now, suppose $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_{k-1}, \boldsymbol{\tau}_k, \ldots, \boldsymbol{\tau}_K$ are linearly independent. Then, there exist unique $\alpha_1, \ldots, \alpha_K$ such that $\boldsymbol{\tau}^* = \sum_{i=1}^{k-1} \alpha_i \boldsymbol{\eta}_i + \sum_{i=k}^K \alpha_i \boldsymbol{\tau}_i$. Further, since $\boldsymbol{\eta}_k \in \text{conv}(\boldsymbol{\tau}_j, \boldsymbol{\tau}^*)^\circ$, there exists $\beta \in (0, 1)$ such that $\boldsymbol{\eta}_k = \beta(\sum_{i=1}^{k-1} \alpha_i \boldsymbol{\eta}_i + \sum_{i=k}^K \alpha_i \boldsymbol{\tau}_i) + (1 - \beta) \boldsymbol{\tau}_k$. By statement 1 of Lemma 3, it suffices to show that $\beta \alpha_k + (1 - \beta) \neq 0$. Suppose to the contrary that $\beta \alpha_k + (1 - \beta) = 0$. Then, $\alpha_k = 1 - \frac{1}{\beta} < 0$. So, it is enough to show that $\boldsymbol{\tau}^* \in \text{conv}(\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_k, \boldsymbol{\tau}_{k+1}, \ldots, \boldsymbol{\tau}_K)$.

   We prove the claim inductively. By Statement 1 of 3, $\boldsymbol{\eta}_1, \boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_K$ are linearly independent, so there exists unique $c_1, \ldots, c_K$ such that $\boldsymbol{\tau}^* = c_1 \boldsymbol{\eta}_1 + \sum_{i=2}^K c_i \boldsymbol{\tau}_i$. Since $\boldsymbol{\eta}_1 \in \text{conv}(\boldsymbol{\tau}_1, \boldsymbol{\tau}^*)$, there exists $\gamma \in (0, 1)$ such that $c_1(\gamma \boldsymbol{\tau}_1 + (1 - \gamma) \boldsymbol{\tau}^*) + \sum_{i=2}^K c_i \boldsymbol{\tau}_i = \boldsymbol{\tau}^*$. If $c_1 < 0$, then $\boldsymbol{\tau}^* \notin \text{conv}(\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_K)$, which is a contradiction. The claim follows inductively.

   Having established the claim, the result follows.

□

Lemma 4 establishes that the operation performed in the last for loop at the end of the Demix algorithm gives the desired distribution.

Lemma 4. *Let $P_1, \ldots, P_L$ be jointly irreducible, $(Q_1, \ldots, Q_{K-1})$ be a permutation of $(P_{i_1}, \ldots, P_{i_{K-1}})$ and $Q_K^1 \in \text{conv}(P_{i_1}, \ldots, P_{i_K})^\circ$. Define the sequence*

$$Q_K^i \longleftarrow \textit{Residue}(Q_K^{i-1} \,|\, Q_{i-1})$$

Then, $Q_K^K = P_{i_K}$.

Lemma 4. Relabel the distributions so that $Q_j = P_j$. Let $\boldsymbol{\mu}_1$ denote the mixture proportion of Q_K^1 and $\boldsymbol{e}_j$ the mixture proportion of P_j. Consider the sequence given by

$$\boldsymbol{\mu}_i \longleftarrow \text{Residue}(\boldsymbol{\mu}_{i-1} \,|\, \boldsymbol{e}_{i-1})$$

We claim that $\boldsymbol{\mu}_i, \boldsymbol{e}_1, \ldots, \boldsymbol{e}_{K-1}$ are linearly independent and $\boldsymbol{\mu}_i$ is the mixture proportion of $\boldsymbol{Q}_K^i$. We prove it inductively. The base case is clear by the hypothesis. Suppose that $\boldsymbol{\mu}_n, \boldsymbol{e}_1, \ldots, \boldsymbol{e}_{K-1}$ are linearly independent and $\boldsymbol{\mu}_n$ is the mixture proportion of Q_K^n. Since $P_1, \ldots, P_L$ are jointly irreducible, by Proposition 4, Residue$(\boldsymbol{\mu}_n \,|\, \boldsymbol{e}_n)$ gives the mixture proportion of Q_2^{n+1}. By statement *1* of Lemma 3, for any $\alpha \geqslant 1$ $\alpha\boldsymbol{\mu}_n + (1-\alpha)\boldsymbol{e}_n, \boldsymbol{e}_1, \ldots, \boldsymbol{e}_{K-1}$ is linearly independent. Therefore, in particular, by statement 1 of Lemma 1, Residue$(\boldsymbol{\mu}_n \,|\, \boldsymbol{e}_n), \boldsymbol{e}_1, \ldots, \boldsymbol{e}_{K-1}$ are linearly independent. The claim follows inductively.

It is enough to show that $\boldsymbol{\mu}_K = \boldsymbol{e}_K$. We claim that for $n \leqslant K-1$, $\boldsymbol{\mu}_n \in \text{conv}(\boldsymbol{e}_n, \ldots, \boldsymbol{e}_K)^\circ$. Consider the base case. $\boldsymbol{\mu}_2 \in \partial\,\text{conv}(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_K)$ by Statement *2* of Lemma 1. Since $\boldsymbol{\mu}_2 = \alpha\boldsymbol{\mu}_1 + (1-\alpha)\boldsymbol{e}_1$ for some $\alpha \geqslant 1$, $\boldsymbol{\mu}_2 \in \text{conv}(\boldsymbol{e}_2, \ldots, \boldsymbol{e}_K)$. Suppose that $\boldsymbol{\mu}_2 \in \partial\,\text{conv}(\boldsymbol{e}_2, \ldots, \boldsymbol{e}_K)$. Without loss of generality, suppose that $\boldsymbol{\mu}_2 \in \text{conv}(\boldsymbol{e}_2, \ldots, \boldsymbol{e}_{K-1})$. Then, since $\boldsymbol{\mu}_2 = \alpha\boldsymbol{\mu}_1 + (1-\alpha)\boldsymbol{e}_1$ for some $\alpha \geqslant 1$, $\boldsymbol{\mu}_1 \in \text{conv}(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_{K-1})$. This implies that $\boldsymbol{\mu}_1 \in \partial\,\text{conv}(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_K)$, which is a contradiction. Hence, $\boldsymbol{\mu}_2 \in \text{conv}(\boldsymbol{e}_2, \ldots, \boldsymbol{e}_K)^\circ$. Inductively, repeating the same argument gives the claim. Therefore, $\boldsymbol{\mu}_{K-1} \in \text{conv}(\boldsymbol{e}_{K-1}, \boldsymbol{e}_K)^\circ$. By considering statement *1* of Lemma 1, it is easy to see that $\boldsymbol{\mu}_K = \boldsymbol{e}_K$. This completes the proof. □

8.5.2 The FaceTest Algorithm

We turn to a key subroutine in the Demix algorithm: the FaceTest algorithm (see Algorithm 6). This algorithm tests whether a set of distributions $Q_1, \ldots, Q_K$ are on the interior of the same face. It computes $R_{i,j} \longleftarrow \text{Residue}(Q_i \,|\, Q_j)$ for every pair; it then computes $\kappa^*(Q_i \,|\, R_{i,j})$ determining whether the two distributions are equal. If there is a single pair for which the two distributions are equal, then $Q_1, \ldots, Q_K$ are not on the interior of the same face. Otherwise, they are on the interior of the same face.

Algorithm 6 FaceTest$(Q_1, \ldots, Q_K)$

```
for i = 1, ..., K do
    for j = 1, ..., K do
        if i ≠ j then
            R_{i,j} ⟵ Residue(Q_i | Q_j)
            if κ*(Q_i | R_{i,j}) = 1 then
                return 0
            end if
        end if
    end for
end for
return 1
```

Proposition 8. *Let $P_1, \ldots, P_K$ be jointly irreducible and $Q_1, \ldots, Q_K \in \text{conv}(P_1, \ldots, P_K)$ be distinct. FaceTest($Q_1, \ldots, Q_K$) returns 1 if and only if $Q_1, \ldots, Q_K$ lie on the interior of the same face of Δ_K.*

Proposition 8. Let $Q_j = \boldsymbol{\eta}_j^T \boldsymbol{P}$ for $\boldsymbol{\eta}_j \in \Delta_K$ and all $j \in [K]$. Suppose that $Q_1, \ldots, Q_K$ lie on the interior of the same face, i.e., there exists $i \in [K]$ such that $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_K \in \text{conv}(\{\boldsymbol{e}_j : j \neq i\})^\circ$. Then, $\mathcal{S}(Q_1) = \ldots = \mathcal{S}(Q_K)$. By statement 3 of Lemma 1, $R_{i,j} \neq Q_i$ for all $i, j \in [K]$ such that $i \neq j$. Then, by Proposition 3, $\kappa^*(Q_i \,|\, R_{i,j}) < 1$. Hence, FaceTest($Q_1, \ldots, Q_K$) returns 1.

Suppose that $Q_1, \ldots, Q_K$ do not all lie on the interior of the same face. Then, there exists Q_i, Q_j that do not lie on the interior of the same face. Without loss of generality, suppose that $\mathcal{S}(Q_j) \nsubseteq S(Q_i)$. Then, by statement 3 of Lemma 1, $R_{i,j} = Q_i$. By Proposition 3, $\kappa^*(R_{i,j} \,|\, Q_i) = 1$. Hence, FaceTest($Q_1, \ldots, Q_K$) returns 0. □

8.5.3 The Demix Algorithm

Theorem 2. We use the notation from the description of the algorithm only replacing K with L and S_i with $\tilde{P}_i$. We prove the result by induction. Suppose $L = 2$. Then, $\boldsymbol{\pi}_i \in \text{conv}(P_1, P_2)$ and $\boldsymbol{\pi}_2 \in \text{conv}(P_1, P_2)$ and $\boldsymbol{\pi}_1 \neq \boldsymbol{\pi}_2$ by linear independence of $\boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$. By Lemma 2 of Blanchard and Scott (2014), it follows that

$$\{\boldsymbol{e}_1, \boldsymbol{e}_2\} = \{\boldsymbol{\mu} : \boldsymbol{\mu} \text{ is the residue of } \boldsymbol{\pi}_i \text{ with respect to } \boldsymbol{\pi}_j, i, j \in [2], i \neq j\}$$

This fact and Proposition 4 imply that if Q_1 is the residue of $\tilde{P}_1$ with respect to $\tilde{P}_2$ and Q_2 is the residue of $\tilde{P}_2$ with respect to $\tilde{P}_1$, then (Q_1, Q_2) is a permutation of (P_1, P_2), completing the base case.

Suppose $L > 2$. With probability 1, $Q \in \text{conv}(\tilde{P}_2, \ldots, \tilde{P}_L)^\circ$. We can write $Q = \boldsymbol{\eta}^T \boldsymbol{P}$ where clearly we have that $\boldsymbol{\eta}$ is a uniformly distributed random vector in $\text{conv}(\boldsymbol{\pi}_2, \ldots, \boldsymbol{\pi}_L)$. Let R be the residue of Q with respect to $\tilde{P}_1$. By Proposition 4, we can write $R = \boldsymbol{\lambda}^T \boldsymbol{P}$ where $\boldsymbol{\lambda}$ is the residue of $\boldsymbol{\eta}$ with respect to $\boldsymbol{\pi}_1$. By statement *2* of Lemma 1, $\boldsymbol{\lambda} \in \partial\Delta_L$.

Step 1: We claim that with probability 1, there is $l \in [L]$ such that $\boldsymbol{\lambda} \in \text{conv}(\{\boldsymbol{e}_j : j \neq l\})^\circ$. Let $B_{i,j} = \text{conv}(\{\boldsymbol{\pi}_1\} \cup \{\boldsymbol{e}_k : k \neq i, j\})$ where $i, j \in [L]$ and $i \neq j$ and let $C = \text{conv}(\boldsymbol{\pi}_2, \ldots, \boldsymbol{\pi}_L)$. First, we argue that $C \cap B_{i,j}$ has affine dimension at most $L-3$.[3] Since $\boldsymbol{\pi}_2, \ldots, \boldsymbol{\pi}_L$ are linearly independent, C has affine dimension $L-2$. Since $\{\boldsymbol{e}_k : k \neq i, j\}$ are linearly independent, $B_{i,j}$ has affine dimension $L-2$ or $L-3$. If $B_{i,j}$ has affine dimension $L-3$, then $C \cap B_{i,j}$ has affine dimension at most $L-3$. So, suppose that $B_{i,j}$ has affine dimension $L-2$. If $C \cap B_{i,j}$ has affine dimension $L-2$, then $\text{aff}\, C = \text{aff}\, B_{i,j}$. Then, in particular, $\boldsymbol{\pi}_1 \in \text{aff}\, C$. But, this contradicts the linear independence of $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_L$. Therefore, $C \cap B_{i,j}$ has affine dimension at most $L-3$.

Because C has affine dimension $L-2$ and $\boldsymbol{\eta}$ is a uniformly distributed random vector in C, with probability 1, $\boldsymbol{\eta} \notin \cup_{i,j \in [L], i \neq j} B_{i,j}$. Since $\boldsymbol{\pi}_1 \in B_{i,j}$ for all i, j and $\boldsymbol{\eta} \in \text{conv}(\boldsymbol{\lambda}, \boldsymbol{\pi}_1)$ by definition, the convexity of $B_{i,j}$ implies that $\boldsymbol{\lambda} \notin \cup_{i \neq j} B_{i,j}$. Since $\boldsymbol{\lambda} \in \partial\Delta_L$, the claim follows.

Step 2: Let $R_i^{(n)}$ be the residue of $m_{\frac{n-1}{n}}(\tilde{P}_i, Q)$ with respect to $\tilde{P}_1$. We claim that there is some finite integer $N \geqslant 2$ such that for all $n \geqslant N$,

$$\text{FaceTest}(R_2^{(n)}, \ldots, R_L^{(n)})$$

returns 1. Let $m_{\frac{n-1}{n}}(\tilde{P}_i, Q) = \boldsymbol{\tau}_n^{(i)^T} \boldsymbol{P}$ for $i \in [L] \backslash \{1\}$; note that $\boldsymbol{\tau}_n^{(i)} = \frac{1}{n}\boldsymbol{\pi}_i + \frac{n-1}{n}\boldsymbol{\eta}$ and, consequently, $\boldsymbol{\tau}_n^{(i)} \in \Delta_L$. Since $\boldsymbol{\eta} \in \text{conv}(\boldsymbol{\pi}_2, \ldots, \boldsymbol{\pi}_L)^\circ$ with probability 1 and $\boldsymbol{\tau}_n^{(i)} \in \text{conv}(\boldsymbol{\pi_i}, \boldsymbol{\eta})^\circ$ for

[3] Note that if $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in \mathbb{R}^L$ are linearly independent and $n \leqslant L$, then $\text{aff}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n)$ has affine dimension $n-1$.

$i \in [L]\backslash\{1\}$ for all $n \in \mathbb{N}$ and $\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_L$ are linearly independent, it follows that for all $n \in \mathbb{N}$ with probability 1, $\boldsymbol{\pi}_1, \boldsymbol{\tau}_n^{(2)}, \ldots, \boldsymbol{\tau}_n^{(L)}$ are linearly independent by statement *3* in Lemma 3. Fix $i \in [L]\backslash\{1\}$. It suffices to show that there is large enough N such that for $n \geqslant N$, Residue$(m\frac{n-1}{n}(\tilde{P}_i, Q) \,|\, \tilde{P}_1)$ is on the same face as R (the other cases are similar). Let $R_i^{(n)} = \boldsymbol{\mu}_i^{(n)^T}\boldsymbol{P}$; by Proposition 4, $\boldsymbol{\mu}_i^{(n)} \in \Delta_L$ and $\boldsymbol{\mu}_i^{(n)}$ is the residue of $\boldsymbol{\tau}_n^{(n)}$ with respect to $\boldsymbol{\pi}_1$. By Proposition 4, it suffices to show that $\mathcal{S}(\boldsymbol{\mu}_i^{(n)}) = \mathcal{S}(\boldsymbol{\lambda})$. As $n \longrightarrow \infty$, $\boldsymbol{\tau}_n^{(i)} = (1 - \frac{n-1}{n})\boldsymbol{\pi}_i + \frac{n-1}{n}\boldsymbol{\eta} \longrightarrow \boldsymbol{\eta}$, hence by statement *2* in Lemma 2, $\left\|\boldsymbol{\mu}_i^{(n)} - \boldsymbol{\lambda}\right\| \longrightarrow 0$. Since with probability 1, $\boldsymbol{\lambda} \in \text{conv}(\{\boldsymbol{e}_j : j \neq l\})^\circ$, it follows that for some large enough n, $\boldsymbol{\mu}_i^{(n)} \in \text{conv}(\{\boldsymbol{e}_j : j \neq l\})^\circ$.

Step 3: The algorithm recurses on $R_2^{(n)}, \ldots, R_L^{(n)}$. Since $\boldsymbol{\pi}_1, \boldsymbol{\tau}_n^{(2)}, \ldots, \boldsymbol{\tau}_n^{(L)}$ are linearly independent, it follows by statement *2* in Lemma 3 that $\boldsymbol{\pi}_1, \boldsymbol{\mu}_2^{(n)}, \ldots, \boldsymbol{\mu}_L^{(n)}$ are linearly independent. Therefore, by the inductive hypothesis, if $(Q_1, \ldots, Q_{L-1}) \longleftarrow$ Demix$(R_2^{(n)}, \ldots, R_L^{(n)})$, then $|\{P_1, \ldots, P_L\}\backslash\{Q_1, \ldots, Q_{L-1}\}| = 1$. Suppose without loss of generality that $(Q_1, \ldots, Q_{L-1})$ is a permutation of $(P_1, \ldots, P_{L-1})$. Note that $\frac{1}{L}\sum_{i=1}^{L}\tilde{P}_i \in \text{conv}(P_1, \ldots, P_L)^\circ$ since $\boldsymbol{\Pi}$ has full rank. Since lines 18-20 generates the same sequence of distributions as the sequence given by Lemma 4 and the conditions of Lemma 4 are satisfied, at the end of line the given for-loop, $Q_L = P_L$ by Lemma 4. The result follows.

□

Remark 1. *We could replace lines (18-22) of the Demix algorithm with a single application of the multi-sample version of κ^* from Blanchard and Scott (2014).*

8.5.4 The Non-Square Demix Algorithm

Now, we examine the non-square case of the demixing problem ($M > L$). Note that knowledge of L is essential to this approach since one must resample exactly L distributions in order to run the square Demix algorithm.

Algorithm 7 NonSquareDemix$(\tilde{P}_1, \ldots, \tilde{P}_M)$

1: $R_1, \ldots, R_L \longleftarrow$ independently uniformly distributed elements in conv$(\tilde{P}_1, \ldots, \tilde{P}_M)$
2: $(Q_1, \ldots, Q_L) \longleftarrow$ Demix$(R_1, \ldots, R_L)$
3: **return** $(Q_1, \ldots, Q_L)$

Corollary 1. *Suppose $M > L$. Let $P_1, \ldots, P_L$ be jointly irreducible and $\boldsymbol{\Pi}$ satisfy **(B′)**. Then, with probability 1, NonSquareDemix$(\tilde{P}_1, \ldots, \tilde{P}_M)$ terminates and returns $(Q_1, \ldots, Q_L)$ such that $(Q_1, \ldots, Q_L)$ is a permutation of $(P_1, \ldots, P_L)$.*

Corollary 1. We can write $R_i = \boldsymbol{\tau}_i^T\boldsymbol{P}$ where $\boldsymbol{\tau}_i \in \Delta_L$ and $i = 1, \ldots, L$. $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_L$ are drawn uniformly independently from a set with positive $(L-1)$-dimensional Lebesgue measure since $\boldsymbol{\Pi}$ has full rank by hypothesis. By Lemma 7, $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_L$ are linearly independent with probability 1. Then, by Theorem 2, with probability 1, Demix$(R_1, \ldots, R_L)$ terminates and Demix$(R_1, \ldots, R_L)$ returns a permutation of $(P_1, \ldots, P_L)$. □

8.6 Decontamination of the Partial Label Model for Classification

8.6.1 Lemmas

Lemma 5 gives a condition on the mixture proportions under which the multi-sample residue is unique. Lemma 2 in Blanchard and Scott (2014) is very similar and is proved in a very similar way. We give a useful generalization here that reproduces many of the same details.

Lemma 5. *Let $l, k \in [L]$. Let $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_L \in \Delta_L$ be linearly independent. We have that condition 1 implies condition 2 and condition 2 implies condition 3.*

1. *There exists a decomposition*

$$\boldsymbol{\tau}_l = \kappa \boldsymbol{e}_k + (1-\kappa)\boldsymbol{\tau}_l'$$

where $\kappa > 0$ and $\boldsymbol{\tau}_l' \in \text{conv}(\{\boldsymbol{\tau}_j : j \neq l\})$. Further, for every $\boldsymbol{e}_i$ such that $i \neq k$, there exists a decomposition:

$$\boldsymbol{e}_i = \sum_{j=1}^{L} a_j \boldsymbol{\tau}_j$$

such that $a_l < \frac{1}{\kappa}$.

2. *Let*

$$\boldsymbol{T} = \begin{pmatrix} \boldsymbol{\tau}_1^T \\ \vdots \\ \boldsymbol{\tau}_L^T \end{pmatrix}$$

The matrix $\boldsymbol{T}$ is invertible and $\boldsymbol{T}^{-1}$ is such that $(\boldsymbol{T}^{-1})_{l,k} > 0$ and $(\boldsymbol{T}^{-1})_{l,i} \leqslant 0$ for $i \neq k$ and $(\boldsymbol{T}^{-1})_{l,k} > (\boldsymbol{T}^{-1})_{j,k}$ for $j \neq l$. In words, the (l,k)th entry in $\boldsymbol{T}^{-1}$ is positive, every other entry in the lth row of $\boldsymbol{T}^{-1}$ is nonpositive and every other entry in the kth column of $\boldsymbol{T}^{-1}$ is strictly less than the (l,k)th entry. [4]

3. *The residue of $\boldsymbol{\tau}_l$ with respect to $\{\boldsymbol{\tau}_j, j \neq l\}$ is $\boldsymbol{e}_k$.*

lemma 5. Without loss of generality, let $l = 1$ and $k = 1$. First, we show that condition *1* implies condition *2*. Suppose that condition *1* holds. Then, there exists $\kappa > 0$ such that

$$\boldsymbol{\tau}_1 = \kappa \boldsymbol{e}_1 + (1-\kappa)\sum_{i=2}^{L} \mu_i \boldsymbol{\tau}_i$$

with $\mu_i \geqslant 0$ for $i \in [L] \backslash \{1\}$. Then,

$$\boldsymbol{e}_1 = \frac{1}{\kappa}(\boldsymbol{\tau}_1 - \sum_{i \geqslant 2}(1-\kappa)\mu_i \boldsymbol{\tau}_i)$$

Hence, the first row of $\boldsymbol{T}^{-1}$ is given by $\frac{1}{\kappa}(1, -(1-\kappa)\mu_2, \cdots, -(1-\kappa)\mu_L)$. This shows that the first row is such that $(\boldsymbol{T}^{-1})_{1,1} > 0$ and $(\boldsymbol{T}^{-1})_{1,i} \leqslant 0$ for $i \neq 1$.

[4] $(\boldsymbol{T}^{-1})_{i,j}$ is the $i \times j$ entry in the matrix $\boldsymbol{T}^{-1}$.

Consider $\boldsymbol{e}_i$ such that $i \neq 1$. Then, we have the relation: $\boldsymbol{e}_i = \sum_{j=1}^{L} a_j \boldsymbol{\tau}_j$, which gives the ith row of $\boldsymbol{T}^{-1}$. By assumption, $a_1 < \frac{1}{\kappa}$, so the $(i,1)$th entry is strictly less than the $(1,1)$th entry. Hence, *2* follows.

Now, we prove that condition *2* implies condition *3*. Suppose condition *2* is true. Consider the optimization problem:

$$\max_{\boldsymbol{\nu},\boldsymbol{\gamma}} \sum_{i=2}^{L} \nu_i \text{ s.t. } \boldsymbol{\tau}_1 = (1 - \sum_{i \geqslant 2} \nu_i)\boldsymbol{\gamma} + \sum_{i=2}^{L} \nu_i \boldsymbol{\tau}_i$$

over $\boldsymbol{\gamma} \in \Delta_L$ and $\boldsymbol{\nu} = (\nu_2, \cdots, \nu_L) \in C_{L-1} = \{(\nu_2, \cdots, \nu_L) : \nu_i \geqslant 0; \sum_{i=2}^{L} \nu \leqslant 1\}$.

By the same argument given in the proof of Lemma 2 of Blanchard and Scott (2014), this optimization problem is equivalent to the program

$$\max_{\boldsymbol{\gamma} \in \Delta_L} \boldsymbol{e}_1^T (\boldsymbol{T}^T)^{-1} \boldsymbol{\gamma} \text{ s.t. } \boldsymbol{\nu}((\boldsymbol{T}^T)^{-1}\boldsymbol{\gamma}) \in C_{L-1}$$

where $\boldsymbol{\nu}(\boldsymbol{\eta}) = \eta_1^{-1}(-\eta_2, \cdots, -\eta_L)$. The above objective is of the the form $\boldsymbol{a}^T\boldsymbol{\gamma}$ where $\boldsymbol{a}$ is the first column of $\boldsymbol{T}^{-1}$. By assumption, for every $i \in [L]$, $\boldsymbol{T}_{1,1}^{-1} > \boldsymbol{T}_{i,1}^{-1}$. Therefore, the unconstrained maximum over $\boldsymbol{\gamma} \in \Delta_L$ is attained uniquely by $\boldsymbol{\gamma} = \boldsymbol{e}_1$. Notice that $(\boldsymbol{T}^T)^{-1}\boldsymbol{e}_1$ is the first row of $\boldsymbol{T}^{-1}$. Denote this vector $\boldsymbol{b} = (b_1, \cdots, b_L)$. We show that $\boldsymbol{\nu}(\boldsymbol{b}) = b_1^{-1}(-b_2, \cdots, -b_L) \in C_{L-1}$. By assumption, $\boldsymbol{b}$ has its first coordinate positive and the other coordinates are nonpositive. Therefore, all of the components of $\boldsymbol{\nu}(\boldsymbol{b})$ are nonnegative. Furthermore, the sum of the components of $\boldsymbol{\nu}(\boldsymbol{b})$ is

$$\sum_{i=2}^{L} \frac{-b_i}{b_1} = 1 - \frac{\sum_{i=1}^{L} b_i}{b_1} = 1 - \frac{1}{b_1} \leqslant 1$$

The last equality follows because the rows of $\boldsymbol{T}^{-1}$ sum to 1 since $\boldsymbol{T}$ is a stochastic matrix. Then, we have $\boldsymbol{\nu}((\boldsymbol{T}^T)^{-1}\boldsymbol{e}_1) \in C_{L-1}$. Consequently, the unique maximum of the optimization problem is attained for $\boldsymbol{\gamma} = \boldsymbol{e}_1$. This establishes *3*. □

Algorithm 8 FindSet($\boldsymbol{B}$)

Input: $\boldsymbol{B}$ is a $N \times K$ binary matrix

```
v ⟵ B_{1,:}
for k = 2, ..., N do
    v' ⟵ min(v, B_{k,:})
    if sum(v') ⩾ 1 then
        v ⟵ v'
    end if
end for
return v
```

Lemma 6 establishes that if $\boldsymbol{S}$ satisfies **(C)**, then we can find a collection of contaminated distributions $\tilde{P}_1, \ldots, \tilde{P}_L$ that only have a single base distribution in common; this observation is essential for establishing that VertexTest behaves appropriately.

Lemma 6. *Let $\boldsymbol{B}$ be a $N \times K$ binary matrix with at least one nonzero entry. If no two columns of $\boldsymbol{B}$ are identical, then there exists $\{i_1, \ldots, i_n\} \subseteq [N]$ such that $\min(\boldsymbol{B}_{i_1,:}, \ldots, \boldsymbol{B}_{i_n,:})$ contains a single 1.*

Proof. The proof is by contraposition. Suppose that there does not exist $\{i_1, \ldots, i_n\} \subseteq [N]$ such that $\min(\boldsymbol{B}_{i_1,:}, \ldots, \boldsymbol{B}_{i_n,:})$ contains a single 1. Without loss of generality, reorder the rows of $\boldsymbol{B}$ such that a non-zero entry of $\boldsymbol{B}$ appears in $\boldsymbol{B}_{1,:}$. Let $\boldsymbol{v}$ denote the vector returned by FindSet($\boldsymbol{B}$) (see Algorithm 8). Then, $\boldsymbol{v}$ contains at least two 1s. Let $v_i = 1$ and $v_j = 1$ where $i \neq j$. Let $\boldsymbol{w}^{(k)} \longleftarrow \min(\boldsymbol{v}, \boldsymbol{B}_{k,:})$. Suppose that there is some k such that $w_i^{(k)} = 1$ and $w_j^{(k)} = 0$. Then, inspection of the Algorithm 8 shows that $\boldsymbol{v}$ would have to already be such that $v_i = 1$ and $v_j = 0$, which is a contradiction. Hence, for all $k \in [N]$, $w_i^{(k)} = w_j^{(k)}$. Then, $\boldsymbol{B}_{:,i} = \boldsymbol{B}_{:,j}$. □

Lemma 7. *If $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in \Delta_L$ are linearly independent and $\boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_L \in \Delta_L$ are random vectors drawn independently from the $(L-1)$-dimensional Lebesgue measure on a set $A \subset \Delta_L$ with positive $(L-1)$-dimensional Lebesgue measure, then $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k, \boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_L$ are linearly independent with probability 1.*

Proof. Suppose that we sample $\boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_L$ sequentially. Then, $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k, \boldsymbol{w}_{k+1}$ are linearly independent if $\boldsymbol{w}_{k+1} \notin \text{span}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k)$. Since $\text{span}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k) \cap \Delta_L$ is a $(k-1)$-dimensional simplex in Δ_L, it has $(L-1)$-dimensional Lebesgue measure 0. Hence, with probability 1, $\boldsymbol{w}_{k+1} \notin \text{span}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k)$. It follows inductively that with probability 1, $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k, \boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_L$ are linearly independent. □

8.6.2 VertexTest Algorithm

To develop some intuition regarding the behavior of VertexTest, consider the following simple example. Suppose that

$$\boldsymbol{S} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}$$

and $Q_1 = P_2$, $Q_2 = P_3$, and $Q_3 = P_1$. Let $\boldsymbol{C}^{(k)}$ denote the value of $\boldsymbol{C}$ in the kth iteration of the loop starting on line 16. Then,

$$\begin{aligned} &\kappa^*(\tilde{P}_1 \,|Q_1) \in (0,1], \quad \kappa^*(\tilde{P}_2 \,|Q_1) = 0, && \kappa^*(\tilde{P}_3 \,|Q_1) \in (0,1] \\ &\kappa^*(\tilde{P}_1 \,|Q_2) = 0, \quad \kappa^*(\tilde{P}_2 \,|Q_2) \in (0,1], && \kappa^*(\tilde{P}_3 \,|Q_2) \in (0,1] \\ &\kappa^*(\tilde{P}_1 \,|Q_3) \in (0,1], \quad \kappa^*(\tilde{P}_2 \,|Q_3) \in (0,1], && \kappa^*(\tilde{P}_3 \,|Q_3) \in (0,1] \end{aligned}$$

$$\boldsymbol{C}^{(0)} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \qquad \boldsymbol{C}^{(1)} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$$

Once we obtain $\boldsymbol{C}^{(0)}$, we know that $Q_3 = P_1$ and, therefore, $Q_2 \neq P_1$ and $Q_1 \neq P_1$. Using this information, we obtain ${\boldsymbol{C}^{(1)}}^T$, which is the desired permutation matrix. On the other hand, if $Q_1 = P_2$, $Q_2 = P_3$, and $Q_3 = \frac{1}{2}P_1 + \frac{1}{2}P_3$, then $\kappa^*(Q_3 \,|\, Q_2) \in (0,1)$ and, therefore, VertexCover would return $(0, I)$.

Lemma 8 establishes the VertexTest algorithm has the desired behavior.

Lemma 8. *Let $\boldsymbol{\eta}_1, \ldots, \boldsymbol{\eta}_L \in \Delta_L$ and $Q_i = \boldsymbol{\eta}_i^T \boldsymbol{P}$ for $i \in [L]$. Suppose that $P_1, \ldots, P_L$ satisfy **(A′)**, $\boldsymbol{\Pi}$ satisfies **(B′)**, and $\boldsymbol{S}$ satisfies **(C)**. Then, VertexTest$(\boldsymbol{S}, (\tilde{P}_1, \ldots, \tilde{P}_L)^T, (Q_1, \ldots, Q_L)^T)$ returns $(1, \boldsymbol{C})$ if and only if $(Q_1, \ldots, Q_L)^T$ is a permutation of $\boldsymbol{P}$. Further, if VertexTest$(\boldsymbol{S}, (\tilde{P}_1, \ldots, \tilde{P}_L)^T, (Q_1, \ldots, Q_L)^T)$ returns $(1, \boldsymbol{C})$, then $\boldsymbol{C}(Q_1, \ldots, Q_L)^T = \boldsymbol{P}$.*

Algorithm 9 VertexTest$(\boldsymbol{S}, (\tilde{P}_1, \ldots, \tilde{P}_L)^T, (Q_1, \ldots, Q_L)^T)$

1: **for** $i = 1, \ldots, L$ **do**
2: **for** $j = 1, \ldots, L$ **do**
3: **if** $\kappa^*(Q_i \,|\, Q_j) \in (0, 1]$ **then**
4: **return** $(0, \boldsymbol{I})$
5: **end if**
6: **end for**
7: **end for**
8: $\boldsymbol{C} \longleftarrow \begin{pmatrix} 1 & \ldots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \ldots & 1 \end{pmatrix} \in \mathbb{R}^{L \times L}$
9: **for** $i = 1, \ldots, L$ **do**
10: **for** $j = 1, \ldots, L$ **do**
11: **if** $\kappa^*(\tilde{P}_j \,|\, Q_i) \in (0, 1]$ **then**
12: $C_{i,:} \longleftarrow \min(C_{i,:}, S_{j,:})$
13: **end if**
14: **end for**
15: **end for**
16: **for** $k = 1$ to L **do**
17: **for** $i = 1$ to L **do**
18: **if** $\boldsymbol{C}_{i,:} == \boldsymbol{e}_j^T$ for some j **then**
19: $\boldsymbol{C}_{:,j} \longleftarrow \boldsymbol{e}_i$
20: **end if**
21: **end for**
22: **end for**
23: **if** $\boldsymbol{C}$ is a permutation matrix **then**
24: **return** $(1, \boldsymbol{C}^T)$
25: **else**
26: **return** $(0, \boldsymbol{C}^T)$
27: **end if**

Proof. **If:** Suppose that $(Q_1, \ldots, Q_L)^T$ is a permutation of $(P_1, \ldots, P_L)^T$. We prove two loop invariants for the loop starting on line 16, which together imply the if-direction. Let $\boldsymbol{C}^{(k)}$ denote the matrix $\boldsymbol{C}$ in the VertexTest algorithm at line 16 after going through the loop starting on line 16 k times.

Claim 1: Let $B_i^{(k)} = \{P_j : C_{i,j}^{(k)} = 1\}$. We show that $Q_i \in B_i^{(k)}$ for all $k \in [L]$ and all $i \in [L]$. Consider the base case: $k = 0$. Fix some $i \in [L]$. Let $A = \{j : \kappa^*(\tilde{P}_j \,|\, Q_i) \in (0, 1]\}$. By construction, $B_i^{(0)} = \{P_k : k \in \cap_{j \in A} \mathcal{S}(\tilde{P}_j)\}$ since

$$\begin{aligned} P_l \in B_i^{(0)} &\iff C_{i,l}^{(0)} = 1 \\ &\iff S_{j,l} = 1 \, \forall j \text{ s.t. } \kappa^*(\tilde{P}_j \,|\, Q_i) \in (0, 1] \\ &\iff P_l \in \{P_k : k \in \cap_{j \in A} \mathcal{S}(\tilde{P}_j)\} \end{aligned}$$

where we used the fact that $S_{j,l} = 1$ iff $P_l \in \mathcal{S}(\tilde{P}_j)$. By statement 3 of Lemma 1, $\kappa^*(\tilde{P}_j \,|\, Q_i) \in$

$(0, 1]$ implies that $\mathcal{S}(Q_i) \subseteq \mathcal{S}(\tilde{P}_j)$, so that $\mathcal{S}(Q_i) \subseteq \cap_{j \in A} \mathcal{S}(\tilde{P}_j)$. Since Q_i is one of the base distributions, $Q_i \in B_i^{(0)}$.

Now, suppose that $Q_i \in B_i^{(n-1)}$ for all $i \in [L]$; we show that $Q_i \in B_i^{(n)}$ for all $i \in [L]$. Suppose that for all l such that $\boldsymbol{C}_{l,:}^{(n-1)} = \boldsymbol{e}_{j_l}^T$ for some j_l, we had $\boldsymbol{C}_{l,:}^{(n-2)} = \boldsymbol{e}_{j_l}^T$. Then, the algorithm does nothing in the nth iteration, so that $Q_i \in B_i^{(n-1)} = B_i^{(n)}$ for all $i \in [L]$.

Now, suppose that there is some l such that $\boldsymbol{C}_{l,:}^{(n-1)} = \boldsymbol{e}_{j_l}$ for some j_l and $\boldsymbol{C}_{l,:}^{(n-2)} \neq \boldsymbol{e}_{j_l}$. By the inductive hypothesis, $Q_l = P_{j_l}$. Then, for all $k \neq l$, we cannot have that $Q_k = P_{j_l}$ because otherwise $(Q_1, \ldots, Q_L)^T$ would not be a permutation of $(P_1, \ldots, P_L)^T$. Therefore, setting $\boldsymbol{C}_{:,j_l}^{(n-1)} \longleftarrow \boldsymbol{e}_l$ cannot remove Q_i from $B_i^{(n-1)}$ for any $i \in [L]$. It follows that $Q_i \in B_i^{(n)}$ for all $i \in [L]$. The claim follows inductively.

Claim 2: We claim that $\boldsymbol{C}^{(k)}$ has at least k columns that are equal to distinct members of the standard basis $\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_L\}$. We give a proof by induction. $k = 0$ is trivial. We show $k = 1$. Since $\boldsymbol{S}$ has distinct columns and at least one nonzero entry, by Lemma 6, there is a set $\{i_1, \ldots, i_l\} \subset [L]$ such that $\min(\boldsymbol{S}_{i_1,:}, \ldots, \boldsymbol{S}_{i_l,:}) = \boldsymbol{e}_i^T$ for some i. Since by assumption $(Q_1, \ldots, Q_L)^T$ is a permutation of $(P_1, \ldots, P_L)^T$, there is some Q_j such that $Q_j = P_i$. Note that by definition of $\boldsymbol{S}$, $\mathcal{S}(Q_j) = \{i\} \subseteq \mathcal{S}(\tilde{P}_k)$ for all $k \in \{i_1, \ldots, i_l\}$. By statement 3 of Lemma 1, $\kappa^*(\tilde{P}_k \,|\, Q_j) \in (0, 1]$ for all $k \in \{i_1, \ldots, i_l\}$. Note that $vec C_{j,i}^{(0)} \neq 0$ by Claim 1; it follows that $\boldsymbol{C}_{j,:}^{(0)} = \boldsymbol{e}_i^T$. Therefore, in the 1st iteration of the loop starting on line 16, the condition $\boldsymbol{C}_{j,:}^{(0)} = \boldsymbol{e}_i^T$ is satisfied and at least one column of $\boldsymbol{C}^{(0)}$ is converted into one of the members of the standard basis in line 19. This proves the case $k = 1$.

Now, suppose that $n \geqslant k$ of the columns of $\boldsymbol{C}^{(k)}$ are equal to distinct members of the standard basis. If $n > k$, then we are done with the inductive step. Therefore, suppose that $n = k$. Let $i_1, \ldots, i_k$ denote the column indices of the columns in $\boldsymbol{C}^{(k)}$ that are equal to members of the standard basis and let $j_1, \ldots, j_k$ denote the row indices of their corresponding non-zero entries. Let $\boldsymbol{S}'$ denote the matrix obtained by deleting the columns $i_1, \ldots, i_k$ from $\boldsymbol{S}$. Since $\boldsymbol{S}$ does not contain two identical columns, $\boldsymbol{S}'$ does not contain two identical columns. We claim that $\boldsymbol{S}'$ has at least one non-zero entry. Suppose to the contrary that $\boldsymbol{S}'$ has only 0 entries. Then, by definition of $\boldsymbol{S}$, $\{\tilde{P}_1, \ldots, \tilde{P}_L\} \subset \text{conv}(P_{i_1}, \ldots, P_{i_k})$. But, if $k < L$, this contradicts the fact that $\tilde{P}_1, \ldots, \tilde{P}_L$ are linearly independent (which we have since $P_1, \ldots, P_L$ are jointly irreducible and $\boldsymbol{\Pi}$ is full rank).

Therefore, by Lemma 6, there is a set of rows $l_1, \ldots, l_t$ of $\boldsymbol{S}'$ such that $\boldsymbol{e}^T \equiv \min(\boldsymbol{S}'_{l_1,:}, \ldots, \boldsymbol{S}'_{l_t,:})$ is a vector of zeros with a unique 1 in one of its entries. Map $\boldsymbol{e}^T$ to the $\boldsymbol{e}_i^T$ obtained by filling in zeros corresponding to the columns $i_1, \ldots, i_k$ that were previously deleted to obtain $\boldsymbol{S}'$ from $\boldsymbol{S}$.

By assumption, there is some Q_j such that $Q_j = P_i$. Note that $i \notin \{i_1, \ldots, i_k\}$ by our construction of $\boldsymbol{e}_i^T$. Note that $j \notin \{j_1, \ldots, j_k\}$ since otherwise for some $j_l \in \{j_1, \ldots, j_k\}$, we would have $P_i = Q_{j_l} = P_{i_l}$ and $P_i \neq P_{i_l}$, which is a contradiction.

Since $\min(\boldsymbol{S}'_{l_1,:}, \ldots, \boldsymbol{S}'_{l_t,:})$ has a unique 1, $\min(\boldsymbol{S}_{l_1,:}, \ldots, \boldsymbol{S}_{l_t,:})$ can only have 1s in positions $i_1, \ldots, i_k, i$ and has at least a single 1 in position i. Therefore, by definition of $\boldsymbol{S}$ and statement 3 of Lemma 1, $\boldsymbol{C}_{j,:}^{(0)}$ can only have nonzero entries in indices $i, i_1, \ldots, i_k$. By assumption, $C_{j,i_1}^{(k)}, \ldots, C_{j,i_k}^{(k)}$ are 0. $\boldsymbol{C}_{j,i}^{(k)} \neq 0$ since that would violate Claim 1. Therefore, $\boldsymbol{C}_{j,:}^{(k)} = \boldsymbol{e}_i^T$. Then,

in the $(k+1)$th iteration of the loop starting on line 16, in line 19, $\boldsymbol{C}_{:,i}^{(k)}$ is converted into one of the standard basis vectors. This establishes the inductive step.

The above two claims imply that once the loop starting on line 16 has terminated, C is a permutation matrix such that if $\boldsymbol{C}_{:,j} = \boldsymbol{e}_i$, then $Q_i = P_j$. Then, $(Q_1, \ldots, Q_L)\boldsymbol{C} = (P_1, \ldots, P_L)$. Taking the transpose of both sides, the result follows.

Only If: Suppose that $(Q_1, \ldots, Q_L)^T$ is not a permutation of $(P_1, \ldots, P_L)^T$. Suppose to the contrary that VertexTest returns $(1, \boldsymbol{C})$ such that $\boldsymbol{C}$ is a permutation matrix; we derive a contradiction. Since $(Q_1, \ldots, Q_L)^T$ is not a permutation of $(P_1, \ldots, P_L)^T$ and $\boldsymbol{C}^{(L)}$ is a permutation matrix by assumption, there exists $k \in [L]$ and $i \in [L]$ such that $Q_i \notin \text{conv}(B_i^{(k)})$. Fix k to be the smallest k such that there exists $i \in [L]$ such that $Q_i \notin \text{conv}(B_i^{(k)})$.

We claim that $k \neq 0$. Let $A = \{j : \kappa^*(\tilde{P}_j \,|\, Q_i) \in (0, 1]\}$. By construction, before entering the loop on line 16, $B_i^{(0)} = \{P_k : k \in \cap_{j \in A} \mathcal{S}(\tilde{P}_j)\}$. By statement 3 of Lemma 1, $\mathcal{S}(Q_i) \subseteq \cap_{j \in A} \mathcal{S}(\tilde{P}_j)$, so that $Q_i \in \text{conv}(B_i^{(0)})$. Thus, $k \neq 0$.

Let $D = \{j : C_{i,j}^{(k-1)} = 1, C_{i,j}^{(k)} = 0\}$. For each $j \in D$, there exists $l_j \in [L]$ such that $\boldsymbol{C}_{l_j,:}^{(k-1)} = \boldsymbol{e}_j^T$. Then, since k is the smallest integer such that there exists $i \in [L]$ such that $Q_i \notin \text{conv}(B_i^{(k)})$, for each $j \in D$, there is l_j such that $Q_{l_j} = P_j$.

Fix i such that $Q_i \in \text{conv}(B_i^{(k-1)})$ and $Q_i \notin \text{conv}(B_i^{(k)})$. If $Q_i = P_j$ for some $j \in D$, then we are done since $\kappa^*(Q_{l_j} \,|\, Q_i) = 1$, implying that VertexTest would output $(0, I)$–a contradiction. Thus, Q_i is not any of the base distributions. There must be some $E \subset B_i^{(k-1)}$ such that $Q_i \in E^\circ$. Then, we must remove some P_j from E such that $Q_i \notin \text{conv}(E \backslash \{P_j\})$. But, we only remove P_j from E if $j \in D$, which means there is l_j such that $Q_{l_j} = P_j$. Since $\mathcal{S}(Q_{l_j}) \subset \mathcal{S}(Q_i)$, $\kappa^*(Q_i \,|\, Q_{l_j}) > 0$, implying that VertexTest would output $(0, I)$–a contradiction.

□

8.6.3 Proof of Theorem 3

Theorem 3. We adopt the notation from the description of Algorithm 2 with the exception that we make explicit the dependence on k by writing $W_i^{(k)}$ instead of W_i and $\bar{Q}_i^{(k)}$ instead of $\bar{Q}_i$. We show that there is a K such that for all $k \geqslant K$, $(S_1^{(k)}, \ldots, S_L^{(k)})^T$ is a permutation of $(P_1, \ldots, P_L)^T$. Then, the result will follow from Lemma 8.

Let $Q_i = \boldsymbol{\tau}_i^T \boldsymbol{P}$, $\bar{Q}_i^{(k)} = \boldsymbol{\tau}^{(k)^T} \boldsymbol{P}$, and $W_i^{(k)} = \boldsymbol{\gamma}_i^{(k)^T} \boldsymbol{P}$. We prove the first claim inductively. First, we consider the base case: there is large enough k such that $W_1^{(k)} = P_i$ for some $i \in [L]$. We will apply Lemma 5. By Lemma 7, $\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_L$ are linearly independent. Therefore, $\text{aff}(\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L)$ gives a hyperplane with an associated open halfspace $\boldsymbol{H}$ that contains $\boldsymbol{\tau}_1$ and at least one $\boldsymbol{e}_j$. We can pick k large enough such that for all $\boldsymbol{e}_j \in \boldsymbol{H}$, $\lambda_k \equiv \frac{1}{k}\boldsymbol{\tau}_1 + \frac{k-1}{k}\boldsymbol{\tau}_1^{(k)} \in \text{conv}(\boldsymbol{e}_j, \boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L)$. Then, for all $\boldsymbol{e}_j \in \boldsymbol{H}$, there exists $\kappa_j > 0$ such that

$$\boldsymbol{\lambda}_k = \kappa_j \boldsymbol{e}_j + (1 - \kappa_j)\tilde{\tau}_j$$

for some $\tilde{\tau}_j \in \text{conv}(\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L)$. We claim that there is a unique smallest κ_j. Suppose to the contrary that there is $i \neq j$ such that $\kappa_i = \kappa_j$. Then,

$$\boldsymbol{\lambda}_k = \kappa \boldsymbol{e}_i + (1-\kappa)\tilde{\tau}_i$$
$$\boldsymbol{\lambda}_k = \kappa \boldsymbol{e}_j + (1-\kappa)\tilde{\tau}_j$$

where $\kappa = \kappa_i = \kappa_j$. Then, $(1-\kappa)(\tilde{\boldsymbol{\tau}}_j - \tilde{\boldsymbol{\tau}}_i) - \kappa(\boldsymbol{e}_i - \boldsymbol{e}_j) = 0$, from which it follows that $\boldsymbol{e}_i - \boldsymbol{e}_j \in \text{span}(\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L)$. But, by Lemma 7, $\boldsymbol{e}_i - \boldsymbol{e}_j, \boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L$ are linearly independent with probability 1 and, hence, we have a contradiction. Therefore, there is a unique j that minimizes κ_j. Note that for all $\boldsymbol{e}_i \notin \boldsymbol{H}$, if we write $\boldsymbol{e}_i = \sum_{l \geqslant 2} a_l \boldsymbol{\tau}_l + a_1 \boldsymbol{\lambda}_k$, then $a_1 \leqslant 0$. Then, by Lemma 5, $\boldsymbol{e}_j$ is the residue of $\boldsymbol{\lambda}_k$ with respect to $\boldsymbol{\tau}_2, \ldots, \boldsymbol{\tau}_L$. Therefore, by Proposition 4, if $W_1^{(k)} \longleftarrow \text{Residue}(\frac{1}{k}Q_1 + (1-\frac{1}{k})\bar{Q}_1 \,|\, \{Q_j\}_{j>1})$, $W_1^{(k)}$ is one of the base distributions.

The inductive step is similar. Suppose that $W_1^{(k)}, \ldots, W_n^{(k)}$ are distinct base distributions for $k \geqslant K_1$. We show that there exists K_2 such that for $k \geqslant K_2$, $W_1^{(k)}, \ldots, W_{n+1}^{(k)}$ are distinct base distributions. Let $k \geqslant K_1$. By hypothesis, $|\{\boldsymbol{\gamma}_1^{(k)}, \ldots, \boldsymbol{\gamma}_n^{(k)}\}| = n$ and $\{\boldsymbol{\gamma}_1^{(k)}, \ldots, \boldsymbol{\gamma}_n^{(k)}\} \subseteq \{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_L\}$. Without loss of generality, let $\boldsymbol{\gamma}_i^{(k)} = \boldsymbol{e}_i$ for $i \leqslant n$. By Lemma 7, $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n, \boldsymbol{\tau}_{n+1}, \ldots, \boldsymbol{\tau}_L$ are linearly independent with probability 1. The rest of the argument is identical to the base case.

The result follows from applying Lemma 8. □

We remark here that the partial label algorithm is basically identical for the non-square case. The above proofs did not make use of the fact that we assumed $M = L$.

8.7 Estimation

8.7.1 ResidueHat Results

Let $A_1, A_2, \ldots$ denote positive constants whose values change from line to line. We introduce the following definitions.

Definition 8. *We say that the distribution F satisfies the* support condition ***(SC)*** *with respect to H if there exists a distribution G and $\gamma \in [0,1]$ such that* $\text{supp}(H) \nsubseteq \text{supp}(G)$ *and $F = (1-\gamma)G + \gamma H$.*

Definition 9. *If*

$$\sup_{S \in \mathcal{S}} |\widehat{F}(S) - F(S)| \xrightarrow{i.p.} 0$$

as $\boldsymbol{n} \longrightarrow \infty$, we say that $\widehat{F} \longrightarrow F$ uniformly (or $\widehat{F}$ converges uniformly to F).

Definition 10. *Let $\widehat{F}$ be a ResidueHat estimator of a distribution F. We say that $\widehat{F}$ satisfies a Uniform Deviation Inequality* ***(UDI)*** *with respect to $\mathcal{S}$ if there exists constants $A_1, A_2 > 0$ such that for large enough* $\min(n_i : i \in \mathcal{D}(\widehat{F}))$*, for all $S \in \mathcal{S}$*

$$|\widehat{F}(S) - F(S)| < A_1 \gamma(\mathcal{D}(\widehat{F}))$$

with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{F})} \frac{1}{n_i}$

Henceforth, for the purposes of abbreviation, we will only say that a ResidueHat estimator satisfies a Uniform Deviation Inequality **(UDI)** and omit "with respect to $\mathcal{S}$" because the context makes this clear.

Definition 11. *Let $\widehat{F}$ and $\widehat{H}$ be ResidueHat estimators. We say that $\widehat{\kappa}(\widehat{F} \mid \widehat{H})$ satisfies a Rate of Convergence **(RC)** if there exists constants $A_1, A_2 > 0$ such that for large enough* $\min(n_i : i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$,

$$|\widehat{\kappa}(\widehat{F} \mid \widehat{H}) - \kappa^*(F \mid H)| \leqslant A_1 \gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$$

with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H})} \frac{1}{n_i}$.

Lemma 9 gives sufficient conditions under which F satisfies **(SC)** with respect to H.

Lemma 9. *Let $P_1, \ldots, P_L$ satisfy **(A″)** and let $F, H \in \text{conv}(P_1, \ldots, P_L)$ such that $F \neq H$. Then, F satisfies **(SC)** with respect to H.*

Proof. Let $A = \arg\min(|B| : B \subseteq \{P_1, \ldots, P_L\}, F, H \in \text{conv}(B))$. Without loss of generality, suppose that $A = \{P_1, \ldots, P_K\}$. F either lies on the boundary of $\text{conv}(P_1, \ldots, P_K)$ or doesn't. If F lies on the boundary of $\text{conv}(P_1, \ldots, P_K)$, then we pick $G = F$ and $\gamma = 0$ to obtain $F = (1-\gamma)F + \gamma H$. Since $P_1, \ldots, P_L$ satisfy **(A″)**, $\text{supp}(H) \nsubseteq \text{supp}(F)$.

Now, suppose that $F \in \text{conv}(P_1, \ldots, P_K)^\circ$. Let $G \longleftarrow \text{Residue}(F \mid H)$; we can write $F = (1-\gamma)G + \gamma H$ for $\gamma \in [0, 1)$ since $F \neq H$. Then, by Statement *2* of Lemma 1 and Proposition 4, G is on the boundary of $\text{conv}(P_1, \ldots, P_K)$.

Without loss of generality, suppose that $G \in \text{conv}(P_1, \ldots, P_{K-1})$. Since $F = (1-\gamma)G + \gamma H \notin \text{conv}(P_1, \ldots, P_{K-1})$, $H \notin \text{conv}(P_1, \ldots, P_{K-1})$. Since $P_1, \ldots, P_L$ satisfy **(A″)**, $\text{supp}(H) \nsubseteq \text{supp}(G)$. This completes the proof. □

Lemma 10 gives sufficient conditions under which an estimator $\widehat{G}$ satisfies a **(UDI)**.

Lemma 10. *Let*

1. *F and H be distributions such that $F \neq H$,*
2. *$G \longleftarrow$ Residue$(F \mid H)$, and*
3. *$\widehat{G} \longleftarrow$ ResidueHat$(\widehat{F} \mid \widehat{H})$.*

*If $\widehat{\kappa}(\widehat{F} \mid \widehat{H})$ satisfies a **(RC)**, $\widehat{H}$ satisfies a **(UDI)**, and $\widehat{F}$ satisfies a **(UDI)**, then $\widehat{G}$ satisfies a **(UDI)**.*

Proof. For the sake of abbreviation, let $\widehat{\kappa} = \widehat{\kappa}(\widehat{F} \mid \widehat{H})$, $\kappa^* = \kappa^*(F \mid H)$, $\widehat{\alpha} = \frac{1}{1-\widehat{\kappa}}$ and $\alpha^* = \frac{1}{1-\kappa^*}$. We claim that there are constants $A_1, A_2 > 0$ such that for sufficiently large $\min(n_i : i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$,

$$\Pr(|\widehat{\alpha} - \alpha^*| < A_1 \gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))) \geqslant 1 - A_2 \sum_{i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F})} \frac{1}{n_i} \tag{2}$$

Since $\widehat{\kappa}$ satisfies a **(RC)**, there exists constants $A_1, A_2 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$,

$$|\widehat{\kappa} - \kappa^*| \leqslant A_1 \gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$$

with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H})} \frac{1}{n_i}$. $\kappa^* < 1$ by Proposition 3, so we can let $\min(n_i : i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$ large enough so that $1 - \kappa^* - A_1 \gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F})) > 0$. Then, on this same event, for large

enough $\min(n_i : i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$,

$$\begin{aligned}|\frac{1}{1-\kappa^*} - \frac{1}{1-\widehat{\kappa}}| &\leqslant A_1 \frac{\gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))}{(1-\kappa^*)(1-\widehat{\kappa})} \\ &\leqslant A_1 \frac{\gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))}{(1-\kappa^*)(1-\kappa^* - A_1\gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F})))}\end{aligned}$$

This proves the claim.

We can write $G = \alpha F + (1-\alpha)H$ with $\alpha \geqslant 1$. Then, by the triangle inequality,

$$\begin{aligned}|\widehat{G} - G| &= |\widehat{\alpha}\widehat{F} + (1-\widehat{\alpha})\widehat{H} - \alpha F - (1-\alpha)H| \\ &\leqslant |\widehat{\alpha}\widehat{F} - \alpha F| + |(1-\widehat{\alpha})\widehat{H} - (1-\alpha)H| \\ &= |\widehat{\alpha}\widehat{F} - \widehat{\alpha}F + \widehat{\alpha}F - \alpha F| + |(1-\widehat{\alpha})\widehat{H} - (1-\widehat{\alpha})H + (1-\widehat{\alpha})H - (1-\alpha)H| \\ &\leqslant |\widehat{\alpha}||\widehat{F} - F| + |\widehat{\alpha} - \alpha| + |1-\widehat{\alpha}||\widehat{H} - H| + |\widehat{\alpha} - \alpha|\end{aligned}$$

Since $\widehat{F}$ satisfies a **(UDI)**, $\widehat{H}$ satisfies a **(UDI)**, and inequality (2) holds, the result follows by the union bound. □

Lemma 11 gives sufficient conditions under which an estimator $\widehat{\kappa}$ satisfies one part of the inequality in a **(RC)**.

Lemma 11. *Let $\widehat{F}$ and $\widehat{H}$ be estimates of distributions F and H that satisfy **(UDI)**s. Then, there exist constants $A_1, A_2 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$,*

$$\kappa^*(F \,|\, H) - \widehat{\kappa}(\widehat{F} \,|\, \widehat{H}) \leqslant A_1\gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$$

with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H})} \frac{1}{n_i}$.

Proof. For abbreviation, let $\kappa^* = \kappa(F \,|\, H)$ and $\widehat{\kappa} = \widehat{\kappa}(\widehat{F} \,|\, \widehat{H})$. Fix $S \in \mathcal{S}$ such that $H(S) > 0$. Since by hypothesis $\widehat{F}$ and $\widehat{H}$ satisfy **(UDI)**, there exist constants $A_1, A_2 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H}))$, with probability at least $1 - A_1[\sum_{i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H})} \frac{1}{n_i}]$,

$$\begin{aligned}|\widehat{F}(S) - F(S)| &< A_2\gamma(\mathcal{D}(\widehat{F})) \\ |\widehat{H}(S) - H(S)| &< A_2\gamma(\mathcal{D}(\widehat{H}))\end{aligned}$$

If $A_2 \leqslant 1$, then $\kappa^* \leqslant \widehat{\kappa}$ on this event since

$$\kappa^* = \inf_{S \in \mathcal{S}} \frac{F(S)}{H(S)} \leqslant \frac{F(S)}{H(S)} \leqslant \frac{\widehat{F}(S) + A_2\gamma(\mathcal{D}(\widehat{F}))}{(\widehat{H}(S) - A_2\gamma(\mathcal{D}(\widehat{H})))_+} \leqslant \frac{\widehat{F}(S) + \gamma(\mathcal{D}(\widehat{F}))}{(\widehat{H}(S) - \gamma(\mathcal{D}(\widehat{H})))_+}$$

and we take the infimimum over $S \in \mathcal{S}$ such that $H(S) > 0$.

Now, suppose that $A_2 > 1$. Let $\min(n_i : i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H}))$ large enough so that $H(S) - 2A_2\gamma(\mathcal{D}(\widehat{H})) > 0$ and with high probability $\widehat{H}(S) - A_2\gamma(\mathcal{D}(\widehat{H}) > 0$. Then, on this event,

$$\begin{aligned}\kappa^* &\leqslant \frac{\widehat{F}(S) + A_2\gamma(\mathcal{D}(\widehat{F}))}{H(S)} \\ &= \frac{\widehat{F}(S) + \gamma(\mathcal{D}(\widehat{F}))}{H(S)} + A_3\gamma(\mathcal{D}(\widehat{F}))\end{aligned}$$

where we fix $A_3 = \frac{(A_2-1)}{H(S)}$. Further,

$$\begin{aligned}
\frac{\widehat{F}(S)+\gamma(\mathcal{D}(\widehat{F}))}{H(S)} &\leqslant \frac{\widehat{F}(S)+\gamma(\mathcal{D}(\widehat{F}))}{\widehat{H}(S)-A_2\gamma(\mathcal{D}(\widehat{H})} \\
&= \frac{\widehat{F}(S)+\gamma(\mathcal{D}(\widehat{F}))}{\widehat{H}(S)-\gamma(\mathcal{D}(\widehat{H}))}\Big[\frac{\widehat{H}(S)-\gamma(\mathcal{D}(\widehat{H}))}{\widehat{H}(S)-A_2\gamma(\mathcal{D}(\widehat{H}))}\Big] \\
&\leqslant \frac{\widehat{F}(S)+\gamma(\mathcal{D}(\widehat{F}))}{\widehat{H}(S)-\gamma(\mathcal{D}(\widehat{H}))}\Big[1+\frac{(A_2-1)\gamma(\mathcal{D}(\widehat{H}))}{\widehat{H}(S)-A_2\gamma(\mathcal{D}(\widehat{H}))}\Big] \\
&\leqslant \frac{\widehat{F}(S)+\gamma(\mathcal{D}(\widehat{F}))}{\widehat{H}(S)-\gamma(\mathcal{D}(\widehat{H}))}\Big[1+\frac{(A_2-1)\gamma(\mathcal{D}(\widehat{H}))}{H(S)-2A_2\gamma(\mathcal{D}(\widehat{H}))}\Big] \\
&\leqslant \frac{\widehat{F}(S)+\gamma(\mathcal{D}(\widehat{F}))}{\widehat{H}(S)-\gamma(\mathcal{D}(\widehat{H}))}\Big[1+\frac{(A_2-1)}{C}\gamma(\mathcal{D}(\widehat{H}))\Big]
\end{aligned}$$

where we pick $0 < C < H(S) - 2A_2\gamma(\mathcal{D}(\widehat{H}))$. Combining the above inequalities and taking the infimum over $S \in \mathcal{S}$ such that $H(S) > 0$ gives

$$\kappa^* \leqslant \widehat{\kappa} + \frac{\widehat{\kappa}(A_2-1)_+}{C}\gamma(\mathcal{D}(\widehat{H}))) + A_3\gamma(\mathcal{D}(\widehat{F}))$$

Noting that there exists $M \in \mathbb{R}$ such that $\widehat{\kappa} < M$ for large enough $\min(n_i : i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H}))$ since there is some $S \in \mathcal{S}$ such that $H(S) > 0$, we obtain the result. □

Lemma 12 gives sufficient conditions under which $\widehat{\kappa}$ satisfies **(RC)**.

Lemma 12. *Let F and H be distributions such that $F \neq H$. If*

- *F satisfies **(SC)** with respect to H,*
- *$\widehat{F}$ satisfies **(UDI)**, and*
- *$\widehat{H}$ satisfies **(UDI)**,*

*then $\widehat{\kappa}(\widehat{F} \,|\, \widehat{H})$ satisfies **(RC)**.*

Proof. By Lemma 11, there exists constants $A_1, A_2 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$,

$$\kappa^*(F \,|\, H) - \widehat{\kappa}(\widehat{F} \,|\, \widehat{H}) \leqslant A_1\gamma(\mathcal{D}(\widehat{H}) \cup \mathcal{D}(\widehat{F}))$$

with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H})} \frac{1}{n_i}$.

The proof of the other direction of the inequality is very similar to the proof of Theorem 2 in Scott (2015). By hypothesis, F satisfies **(SC)** with respect to H, so there exists a distribution G such that $F = \gamma G + (1-\gamma)H$ for some $\gamma \in [0, 1]$ and $\operatorname{supp}(H) \nsubseteq \operatorname{supp}(G)$. Therefore, we have that G is irreducible with respect to H and, by Proposition 3, $\kappa^*(F \,|\, H) = \gamma$. For abbreviation, let $\kappa^* = \kappa^*(F \,|\, H)$ and $\widehat{\kappa} = \widehat{\kappa}(\widehat{F} \,|\, \widehat{H})$. Since $\operatorname{supp}(H) \nsubseteq \operatorname{supp}(G)$, there exists an open set O such that

$$\frac{G(O)}{H(O)} = (1-\gamma)\frac{G(O)}{H(O)} + \gamma = \kappa^*$$

Then, since $\mathcal{S}$ contains a generating set for the standard topology on $\mathbb{R}^d$, there exists $S \in \mathcal{S}$ such that

$$\frac{G(S)}{H(S)} = \kappa^*$$

Since by hypothesis $\widehat{F}$ and $\widehat{H}$ satisfy **(UDI)**, there exist constants $A_3, A_4 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H}))$, with probability at least $1 - A_3[\sum_{i \in \mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H})} \frac{1}{n_i}]$,

$$\begin{aligned}\widehat{\kappa} &\leqslant \frac{F(S) + A_4\gamma(\mathcal{D}(\widehat{F}))}{(H(S) - A_4\gamma(\mathcal{D}(\widehat{H})))_+} \\ &\leqslant \frac{F(S) + \epsilon}{(H(S) - \epsilon)_+}\end{aligned}$$

where $\epsilon = 2A_4\gamma(\mathcal{D}(\widehat{F}) \cup \mathcal{D}(\widehat{H}))$. The rest of the proof is identical to the proof of Theorem 2 from Scott (2015) and, therefore, we omit it. □

The following theorem gives sufficient conditions under which a ResidueHat estimator satisfies **(UDI)**. It is the basis of Proposition 5.

Lemma 13. *If $P_1, \ldots, P_L$ satisfy **(A″)** and $\widehat{G}$ is a ResidueHat estimator of a distribution $G \in \text{conv}(P_1, \ldots, P_L)$, then $\widehat{G}$ satisfies **(UDI)**.*

Proof. A ResidueHat estimator is defined recursively in terms of other ResidueHat estimators, which are empirical distributions in the base case. Let $\widehat{F}_1, \ldots, \widehat{F}_k$ denote the ResidueHat estimators in terms of which $\widehat{G}$ is defined and let $F_1, \ldots, F_k$ denote the distributions that they estimate. Consider the directed graph (V, E) defined as follows: $V = \{v_1, \ldots, v_k, g\}$ where v_i corresponds to estimator $\widehat{F}_i$ and g to estimator $\widehat{G}$. Let $(v_i, v_j) \in E$ if there exists $l \in [k]$ such that either $\widehat{F}_j \longleftarrow \text{ResidueHat}(\widehat{F}_i \,|\, \widehat{F}_l)$ or $\widehat{F}_j \longleftarrow \text{ResidueHat}(\widehat{F}_l \,|\, \widehat{F}_i)$; define edges similarly for the nodes corresponding to the estimators $\widehat{G}$, $\widehat{F}_i$ and $\widehat{F}_j$ such that $\widehat{G} \longleftarrow \text{ResidueHat}(\widehat{F}_i \,|\, \widehat{F}_j)$. Notice that *(i)* (V, E) is a directed acyclic graph, *(ii)* the set of nodes in (V, E) with no incoming edges correspond to the empirical distributions among $\widehat{F}_1, \ldots, \widehat{F}_k$, *(iii)* the only node with no outgoing edges is g, and *(iv)* every node with incoming edges has exactly 2 incoming edges.

We will show that $\widehat{F}_1, \ldots, \widehat{F}_k, \widehat{G}$ all satisfy **(UDI)** by considering the estimators in the following order: first, we consider estimators whose nodes have incoming edges from nodes representing empirical distributions. Then, we only consider a node if its incoming edges come from nodes that represent estimators that have either been shown to satisfy **(UDI)** or are empirical distributions. There are four cases that we must consider. Let $\widehat{F}_l \longleftarrow \text{ResidueHat}(\widehat{F}_i \,|\, \widehat{F}_j)$.

1. $\widehat{F}_i$ and $\widehat{F}_j$ are empirical distributions,
2. $\widehat{F}_i$ is an empirical distribution and $\widehat{F}_j$ satisfies **(UDI)**,
3. $\widehat{F}_i$ satisfies **(UDI)** and $\widehat{F}_j$ is an empirical distribution, and
4. $\widehat{F}_i$ and $\widehat{F}_j$ satisfy **(UDI)**.

Suppose that $\widehat{F}_i$ and $\widehat{F}_j$ are empirical distributions. Since $\widehat{F}_i$ and $\widehat{F}_j$ are empirical distributions, the VC inequality applies to $\widehat{F}_i$ and $\widehat{F}_j$. Therefore, $\widehat{F}_i$ and $\widehat{F}_j$ satisfy **(UDI)**. Since $\widehat{F}_l$ is a ResidueHat estimator, by definition $F_i, F_j \in \text{conv}(P_1, \ldots, P_L)$ and $F_i \neq F_j$. Then, by Lemma 9, F_i satisfies **(SC)** with respect to F_j. Then, by Lemma 12, $\widehat{\kappa}(\widehat{F}_i \,|\, \widehat{F}_j)$ satisfies **(RC)**. Then, all of the assumptions of Lemma 10 are satisfied, so $\widehat{F}_l$ satisfies **(UDI)**. Note that $F_l \in \text{conv}(P_1, \ldots, P_L)$ by Proposition 4. The rest of the cases are similar. Making the appropriate argument at each node in the directed acylcic graph eventually shows that $\widehat{G}$ satisfies **(UDI)**. □

Proposition 5. By Lemma 13, $\widehat{G}$ satisfies **(UDI)**. Explicitly, there exist constants $A_1, A_2 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{G}))$ with probability at least $1 - A_1 \sum_{i \in \mathcal{D}(\widehat{G})} \frac{1}{n_i}$, $\widehat{G}$ satisfies for every $S \in \mathcal{S}$,

$$|\widehat{G}(S) - G(S)| < A_2 \gamma(\mathcal{D}(\widehat{G})) = A_2 \sum_{i \in \mathcal{D}(\widehat{G})} \epsilon_i(\frac{1}{n_i}) \longrightarrow 0$$

as $\boldsymbol{n} \longrightarrow \infty$. □

8.7.2 The FaceTestHat Algorithm

The FaceTestHat algorithm (see Algorithm 10) differs from the FaceTest algorithm in that it requires the specification of $\epsilon > 0$. As will become clear in the following section, this difference introduces further changes in DemixHat.

Algorithm 10 FaceTestHat($\widehat{Q}_1, \cdots, \widehat{Q}_K \,|\, \epsilon$)

```
for i = 1, ··· , K do
    for j = 1, ··· , K do
        if i ≠ j then
            R̂_{i,j} ⟵ ResidueHat(Q̂_i | Q̂_j)
            if κ̂(Q̂_i | R̂_{i,j}) ⩾ 1 − ϵ then
                return 0
            end if
        end if
    end for
end for
return 1
```

Lemma 14. *Let $\epsilon > 0$. Suppose that $P_1, \ldots, P_L$ satisfy **(A″)** and that $Q_1, \ldots, Q_K$ are distinct distributions lying in* $\text{conv}(P_1, \ldots, P_L)$. *Let $\widehat{Q}_i$ be a ResidueHat estimate of Q_i $\forall i \in [K]$. Then, as $\boldsymbol{n} \longrightarrow \infty$, with probability tending to* 1, *if FaceTestHat($\widehat{Q}_1, \cdots, \widehat{Q}_K \,|\, \epsilon$) returns* 1, *then $Q_1, \cdots, Q_K$ are in the interior of the same face.*

Lemma 14. Let $\epsilon > 0$, $\kappa^*_{i,j} = \kappa^*(Q_i \,|\, R_{i,j})$ and $\widehat{\kappa}_{i,j} = \widehat{\kappa}(\widehat{Q}_i \,|\, \widehat{R}_{i,j})$. We prove the contrapositive. Suppose that $Q_1, \ldots, Q_K$ are not on the interior of the same face. Then, by Proposition 8, FaceTest$(Q_1, \ldots, Q_K)$ returns 0, which occurs if and only if there exist $i \neq j$ such that $\kappa^*_{i,j} = 1$.

Since $\widehat{Q}_i$ and $\widehat{Q}_j$ are ResidueHat estimators, by Lemma 13, $\widehat{Q}_i$ and $\widehat{Q}_j$ satisfy **(UDI)**. Since $Q_i \neq Q_j$ by hypothesis, $\widehat{R}_{i,j}$ is a ResidueHat estimator and, therefore, satisfies a **(UDI)** by Lemma 13.

By Lemma 11, there exist constants $A_1, A_2 > 0$ such that there is some N_0 such that if $\min(n_s : s \in \mathcal{D}(\widehat{R}_{i,j})) \geqslant N_0$, with probability at least $1 - A_1 \sum_{s \in \mathcal{D}(\widehat{R}_{i,j})} \frac{1}{n_s}$, we have that $\widehat{\kappa}_{i,j} + A_2\gamma(\mathcal{D}(\widehat{R}_{i,j})) \geqslant \kappa^*_{i,j}$. Let N_1 be such that if $\min(n_s : s \in \mathcal{D}(\widehat{R}_{i,j})) \geqslant N_1$, then $A_2\gamma(\mathcal{D}(\widehat{R}_{i,j})) < \epsilon$. Then, if $\min(n_s : s \in \mathcal{D}(\widehat{R}_{i,j})) \geqslant \max(N_0, N_1)$, with probability at least $1 - A_1 \sum_{s \in \mathcal{D}(\widehat{R}_{i,j})} \frac{1}{n_s}$

$$\begin{aligned}\widehat{\kappa}_{i,j} &\geqslant \kappa^*_{i,j} - \gamma(\mathcal{D}(\widehat{R}_{i,j})) \\ &\geqslant 1 - \epsilon\end{aligned}$$

So, FaceTestHat($\{\widehat{Q}_1, \cdots, \widehat{Q}_K\}$) returns 0. □

Lemma 15. *Suppose that $P_1, \ldots, P_L$ satisfy **(A″)** and that $Q_1, \ldots, Q_K$ are distinct distributions lying on the interior of the same face of* conv$(P_1, \ldots, P_L)$. *Let $\widehat{Q}_i$ be a ResidueHat estimate of Q_i $\forall i \in [K]$. Then, there exists $\delta > 0$ such that if $\delta > \epsilon > 0$, then with probability increasing to 1 as $\boldsymbol{n} \longrightarrow \infty$, FaceTestHat($\widehat{Q}_1, \ldots, \widehat{Q}_L \mid \epsilon$) returns* 1.

Proof. Let $\kappa^*_{i,j} = \kappa^*(Q_i, \mid R_{i,j})$ and $\widehat{\kappa}_{i,j} = \widehat{\kappa}(\widehat{Q}_i \mid \widehat{R}_{i,j})$. By Proposition 8, since $Q_1, \ldots, Q_K$ are distinct distributions lying on the interior of the same face of conv$(P_1, \ldots, P_L)$, $\kappa^*_{i,j} < 1$ for all $i, j \in [K]$ not equal. We verify the conditions of Lemma 12 to obtain consistency of each $\widehat{\kappa}_{i,j}$. Since every $\widehat{Q}_i$ is a ResidueHat estimator, every $\widehat{Q}_i$ satisfies **(UDI)** by Lemma 13. Since $Q_i \neq Q_j$, $\widehat{R}_{i,j}$ is a ResidueHat estimator and, therefore, satisfies **(UDI)** by Lemma 13. Since for all $i \in [K]$, Q_i in the interior of the same face, by Statement *2* of Lemma 1 and Proposition 4, we have that $R_{i,j} \neq Q_i$ for all $i, j \in [K]$. Since $P_1, \ldots, P_L$ satisfy **(A″)**, $R_{i,j}, Q_i \in$ conv$(P_1, \ldots, P_L)$, and $Q_i \neq R_{i,j}$, it follows by Lemma 9 that Q_i satisfies **(SC)** with respect to $R_{i,j}$. Since the conditions of Lemma 12 are satisfied, each $\widehat{\kappa}_{i,j}$ satisfies **(RC)** and is therefore consistent to $\kappa^*_{i,j}$. Pick $\delta = \min_{i,j} \frac{1-\kappa^*_{i,j}}{2}$. Then, the result follows. □

8.7.3 The DemixHat Algorithm

The DemixHat algorithm (see Algorithm 11) differs from the Demix algorithm in that it applies the FaceTestHat algorithm with increasing values of $\epsilon > 0$. At some point, it is guaranteed to estimate distributions on the same face of the simplex and pick ϵ sufficiently small so that FaceTestHat can determine that they are on the same face with probability tending to 1.

Theorem 4. Note that every estimator of a distribution in the DemixHat algorithm is a ResidueHat estimator since the Demix algorithm only considers distributions that are in conv$(P_1, \ldots, P_L)$ and only computes Residue$(F \mid H)$ if $F \neq H$. Therefore, every estimator of a distribution of the DemixHat algorithm satisfies the assumptions of Lemma 13.

First, we argue that DemixHat uses ResidueHat estimators that are recursively defined in terms of a finite number of ResidueHat esimators so that the constants in the uniform deviation inequalities associated with the ResidueHat estimators do not go to infinity. We find a very loose bound. DemixHat calls itself at most $L - 1$ times and in each call recurses on at most $L - 1$ ResidueHat estimators and calculates at most $L - 1$ more ResidueHat estimators. Therefore, each ResidueHat estimator is recursively defined in terms of at most $(L - 1)^3$ ResidueHat estimators.

Second, we argue that with probability increasing to 1 as $\boldsymbol{n} \longrightarrow \infty$, DemixHat eventually terminates. Let A_i denote the event that DemixHat recurses on i distributions lying in the interior of a i-face in the $(L-i)$th recursive call. Consider A_{L-1}. Let $\widehat{R}^k_i$ denote the estimate of the ith distribution in line 12 in the kth

Algorithm 11 DemixHat($\widehat{S}_1, \ldots, \widehat{S}_K$)

Input: $\widehat{S}_1, \ldots, \widehat{S}_K$ are ResidueHat estimates

1: **if** $K = 2$ **then**
2: $\widehat{Q}_1 \longleftarrow$ ResidueHat($\widehat{S}_1 \,|\, \widehat{S}_2$)
3: $\widehat{Q}_2 \longleftarrow$ ResidueHat($\widehat{S}_2 \,|\, \widehat{S}_1$)
4: **return** $(\widehat{Q}_1, \widehat{Q}_2)$
5: **else**
6: $\widehat{Q} \longleftarrow$ uniformly distributed random element from conv($\widehat{S}_2, \ldots, \widehat{S}_K$)
7: $n \longleftarrow 1$
8: $T \longleftarrow 0$
9: **while** $T == 0$ **do**
10: $n \longleftarrow n + 1$
11: **for** $i = 2, \cdots, K$ **do**
12: $\widehat{R}_i \longleftarrow$ ResidueHat($m_{\frac{n-1}{n}}(\widehat{S}_i, \widehat{Q}) \,|\, \widehat{S}_1$)
13: **end for**
14: $\epsilon \longleftarrow \frac{1}{2^{n+1}}$
15: $T \longleftarrow$ FaceTestHat($\widehat{R}_2, \cdots, \widehat{R}_K \,|\, \epsilon$)
16: **end while**
17: $(\widehat{Q}_1, \cdots, \widehat{Q}_{K-1})^T \longleftarrow$ DemixHat($\widehat{R}_2, \cdots, \widehat{R}_K$)
18: $\widehat{Q}_K \longleftarrow \frac{1}{K} \sum_{i=1}^{K} \widehat{S}_i$
19: **for** $i = 1, \ldots, K-1$ **do**
20: $\widehat{Q}_K \longleftarrow$ ResidueHat($\widehat{Q}_K \,|\, \widehat{Q}_i$)
21: **end for**
22: **return** $(\widehat{Q}_1, \cdots, \widehat{Q}_K)^T$
23: **end if**

iteration of the while loop starting on line 9 of the DemixHat algorithm and let R_i^k denote the corresponding distribution. By Theorem 2, there exists a smallest N such that for $n \geqslant N$, every R_i^n lies in the interior of the same face. By Lemma 14, with probability increasing to 1, for any $\tau > 0$, FaceTestHat($\widehat{R}_2^k, \ldots, \widehat{R}_L^k \,|\, \tau$) returns 0 if $k < N$ since $R_2^k, \ldots, R_L^k$ are not on the interior of the same face. By Lemma 15, there exists $\delta > 0$ such that for $\delta > \epsilon > 0$, with probability increasing to 1, for $n \geqslant N$, FaceTestHat($R_2^n, \ldots, R_L^n \,|\, \epsilon$) returns 1. Hence, with probability increasing to 1, the event A_{L-1} occurs. Applying the same argument to A_i for $i < L - 1$ and taking the union bound shows that DemixHat terminates with probability increasing to 1.

Now, we can complete the proof. Under the assumptions of Theorem 2, there is a permutation σ such that for each distribution Q_i estimated by $\widehat{Q}_i$, $P_{\sigma(i)} = Q_i$. By Proposition 5, as $\boldsymbol{n} \longrightarrow \infty$, $\widehat{Q}_i$ converges uniformly to Q_i. The result follows. □

8.7.4 PartialLabelHat Algorithm

We introduce a finite sample algorithm for decontamination of the partial label model. It combines the DemixHat algorithm and an empirical version of the VertexTest algorithm. We make an assumption that simplifies our algorithm: $\boldsymbol{S}$ satisfies

(D) there does not exist $i, j \in [L]$ such that $\boldsymbol{S}_{i,:} = \boldsymbol{e}_j^T$.

In words, this says that there is no contaminated distribution $\tilde{P}_i$ and base distribution P_j such that $\tilde{P}_i = P_j$. We emphasize that we make this assumption only to simplify the presentation and development of the algorithm; one can reduce any instance of a partial label model satisfying **(B′)**, **(C)**, and **(A″)** to an instance of a partial label model that also satisfies **(D)**. We provide a sketch of such a reduction. Let $J = \{i : \boldsymbol{S}_{i,:} = \boldsymbol{e}_j^T \text{ for some } j \in [L]\} = \{j_1, \ldots, j_k\}$, the set of indices of contaminated distributions that are equal to some base distribution. Compute Residue$(\tilde{P}_i \,|\, \tilde{P}_{j_1})$ for $i \in [L] \backslash J$ if there is l such that $\boldsymbol{S}_{i,l} = \boldsymbol{S}_{j_1,l} = 1$. Replace $\tilde{P}_i$ with Residue$(\tilde{P}_i \,|\, \tilde{P}_{j_1})$ (and call it $\tilde{P}_i$ for simplicity of presentation). Update $\boldsymbol{S}$ and remove j_1 from J. Repeat this procedure until J is empty. Then, there will be k $\tilde{P}_i$ lying in a k-face of Δ_L (for some k) that are not equal to any of the base distributions and the other contaminated distributions will be equal to base distributions. Then, it suffices to solve the instance of the partial label model on the k-face, which satisfies **(D)**.

Algorithm 12 PartialLabelHat$(\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T)$

$(\widehat{Q}_1, \ldots, \widehat{Q}_L)^T \longleftarrow$ DemixHat$(\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)$
FoundVertices $\longleftarrow 0$
$k \longleftarrow 2$
while FoundVertices $== 0$ **do**
  (FoundVertices, $\boldsymbol{C}$) $\longleftarrow$ VertexTest$(\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T, (\widehat{Q}_1, \ldots, \widehat{Q}_L)^T, \frac{1}{k})$
  $k \longleftarrow k + 1$
end while
return $\boldsymbol{C}(\widehat{Q}_1, \ldots, \widehat{Q}_L)^T$

Lemma 16. *Suppose that $P_1, \ldots, P_L$ satisfy **(A″)**, $\boldsymbol{\Pi}$ satisfies **(B′)**, and $\boldsymbol{S}$ satisfies **(C)** and **(D)**. Suppose that $\widehat{Q}_1, \ldots, \widehat{Q}_L$ are ResidueHat estimators that are a permutation of $P_1, \ldots, P_L$. Then, there exists $\delta > 0$ such that*

(i) *if $\delta > \epsilon > 0$, then with probability increasing to 1, as $\boldsymbol{n} \longrightarrow \infty$, VertexTest$(\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T, (\widehat{Q}_1, \ldots, \widehat{Q}_L)^T, \epsilon)$ returns a permutation matrix $\boldsymbol{C}$ such that $\forall i$, $\boldsymbol{C}_{i,:}(\widehat{Q}_1, \ldots, \widehat{Q}_L)^T$ is a ResidueHat estimator of P_i;*

(ii) *if $\epsilon > \delta > 0$, then with probability increasing to 1, as $\boldsymbol{n} \longrightarrow \infty$, either VertexTest$(\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T, (\widehat{Q}_1, \ldots, \widehat{Q}_L)^T, \epsilon)$ returns the same permutation matrix $\boldsymbol{C}$ as in* **(i)** *or it returns the value $(0, \boldsymbol{C})$ indicating that it did not find the permutation matrix.*

Proof. **(i)** Fix i, j. Let $\tilde{P}_j$ and Q_i denote the distributions that $\tilde{P}_j^\dagger$ and $\widehat{Q}_i$ estimate. We show that there exists $\delta_{i,j} > 0$ such that if $0 < \epsilon < \delta_{i,j}$, then $\kappa^*(\tilde{P}_j \,|\, Q_i) \in (0, 1]$ if and only if FaceContainHat$(\tilde{P}_j^\dagger, \widehat{Q}_i \,|\, \epsilon)$ returns 1. We will take $\delta = \min_{i,j} \delta_{i,j}$; **(i)** will follow from Lemma 8 and the observation that VertexTestHat$(\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T, (\widehat{Q}_1, \ldots, \widehat{Q}_L)^T, \epsilon)$ and VertexTest$(\boldsymbol{S}, (\tilde{P}_1, \ldots, \tilde{P}_L)^T, (Q_1, \ldots, Q_L)^T)$ have identical behavior if the claim holds.

Let $\widehat{R} \longleftarrow$ ResidueHat$(\tilde{P}_j \,|\, \widehat{Q}_i)$ and R denote the distribution that $\widehat{R}$ estimates. Let $\widehat{\kappa} = \widehat{\kappa}(\tilde{P}_j^\dagger \,|\, \widehat{R})$. We verify the conditions of Lemma 11 to obtain an upper bound on $\widehat{\kappa}$. It follows from the hypothesis that $Q_i \neq \tilde{P}_j$. Further, since $\widehat{Q}_i$ is a ResidueHat estimator by hypothesis, $\tilde{P}_j^\dagger$ is an empirical distribution, and $Q_i, \tilde{P}_j \in \text{conv}(P_1, \ldots, P_L)$, we have that $\widehat{R}$ is a ResidueHat estimator by definition. By Lemma 13, $\widehat{R}$ satisfies **(UDI)**; by the VC inequality, $\tilde{P}_j^\dagger$ satisfies **(UDI)**. Then, by Lemma 11, there exist

Algorithm 13 VertexTestHat($\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T, (\widehat{Q}_1, \ldots, \widehat{Q}_L)^T, \epsilon$)

$\boldsymbol{D} \longleftarrow \begin{pmatrix} 1 & \ldots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \ldots & 1 \end{pmatrix} \in \mathbb{R}^{L \times L}$
for $i = 1, \ldots, L$ **do**
  for $j = 1, \ldots, L$ **do**
    if FaceContainHat($\tilde{P}_j^\dagger, \widehat{Q}_i \,|\epsilon) == 1$ **then**
      $\boldsymbol{D}_{i,:} \longleftarrow \min(\boldsymbol{D}_{i,:}, \boldsymbol{S}_{j,:})$
    end if
  end for
end for
for $k = 1$ to L **do**
  for $i = 1$ to L **do**
    if $\boldsymbol{D}_{i,:} == \boldsymbol{e}_j^T$ for some j **then**
      $\boldsymbol{D}_{:,j} \longleftarrow \boldsymbol{e}_i$
    end if
  end for
end for
if $\boldsymbol{D}$ is a permutation matrix **then**
  return $(1, \boldsymbol{D}^T)$
else
  return $(0, \boldsymbol{D}^T)$
end if

Algorithm 14 FaceContainHat($\widehat{Q}_1, \widehat{Q}_2 \,|\, \epsilon$)

$\widehat{R} \longleftarrow$ ResidueHat($\widehat{Q}_1 \,|\, \widehat{Q}_2$)
if $\widehat{\kappa}(\widehat{Q}_1 \,|\, \widehat{R}) \leqslant 1 - \epsilon$ **then**
  return 1
else
  return 0
end if

constants $A_1, A_2 > 0$ such that for large enough $\min(n_i : i \in \mathcal{D}(\widehat{R}) \cup \mathcal{D}(\tilde{P}_j^\dagger))$,

$$\kappa^*(\tilde{P}_j \,|\, R) - \widehat{\kappa}(\tilde{P}_j^\dagger \,|\, \widehat{R}) \leqslant A_1 \gamma(\mathcal{D}(\widehat{R}) \cup \mathcal{D}(\tilde{P}_j^\dagger)) \tag{3}$$

with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{R}) \cup \mathcal{D}(\tilde{P}_j^\dagger)} \frac{1}{n_i}$.

Now, suppose that $\kappa^*(\tilde{P}_j \,|\, Q_i) = 0$. Then, $\tilde{P}_j = R$ by definition of $\kappa^*(\tilde{P}_j \,|\, Q_i)$. Then, $\kappa^*(\tilde{P}_j \,|\, R) = 1$ by definition of $\kappa^*(\tilde{P}_j \,|\, R)$. Let $\epsilon > 0$ and $\min(n_i : i \in \mathcal{D}(\widehat{R}) \cup \mathcal{D}(\tilde{P}_j^\dagger))$ large enough so that (3) holds with the given probability and $A_1 \gamma(\mathcal{D}(\widehat{R}) \cup \mathcal{D}(\tilde{P}_j^\dagger)) \leqslant \epsilon$. Then, with probability at least $1 - A_2 \sum_{i \in \mathcal{D}(\widehat{R}) \cup \mathcal{D}(\tilde{P}_j^\dagger)} \frac{1}{n_i}$,

$$\widehat{\kappa}(\tilde{P}_j^\dagger \,|\, \widehat{R}) \geqslant \kappa^*(\tilde{P}_j \,|\, R) - A_1 \gamma(\mathcal{D}(\widehat{R}) \cup \mathcal{D}(\tilde{P}_j^\dagger)) \geqslant 1 - \epsilon \tag{4}$$

Therefore, for any $\epsilon > 0$, with probability tending to 1 as $\boldsymbol{n} \longrightarrow \infty$, FaceContainHat$(\tilde{P}_j^\dagger, \widehat{Q}_i \,|\epsilon)$ returns 0.

Suppose that $\kappa^*(\tilde{P}_j \,|\, Q_i) \in (0, 1]$. Since $\tilde{P}_j \neq Q_i$ by hypothesis, $\kappa^*(\tilde{P}_j \,|\, Q_i) \in (0, 1)$. Therefore, $R \neq \tilde{P}_j$ so that $\kappa^*(\tilde{P}_j \,|\, R) < 1$. Since $P_1, \ldots, P_L$ satisfy **(A″)**, by Lemma 9, $\tilde{P}_j$ satisfies **(SC)** with respect to R. Therefore, by Lemma 12, $\widehat{\kappa}(\tilde{P}_j^\dagger \,|\, \widehat{R})$ satisfies **(RC)**. In particular, $\widehat{\kappa}(\tilde{P}_j^\dagger \,|\, \widehat{R})$ is consistent to $\kappa^*(\tilde{P}_j \,|\, R)$.

Let $\delta_{i,j} = \frac{1-\kappa^*(\tilde{P}_j \,|\, R)}{2}$. Then, for $\epsilon < \delta_{i,j}$, with probability tending to 1, as $\boldsymbol{n} \longrightarrow \infty$, $\widehat{\kappa} < 1 - \epsilon$. Therefore, FaceContainHat$(\tilde{P}_j^\dagger, \widehat{Q}_i \,|\, \epsilon)$ returns 1.

(ii) Let $\delta = \min_{i,j} \delta_{i,j}$ and suppose that $\epsilon > \delta > 0$. We again consider compare the execution of

$$\text{VertexTestHat}(\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T, (\widehat{Q}_1, \ldots, \widehat{Q}_L)^T, \epsilon)$$
$$\text{VertexTest}(\boldsymbol{S}, (\tilde{P}_1, \ldots, \tilde{P}_L)^T, (Q_1, \ldots, Q_L)^T)$$

Let $\boldsymbol{D}^{(k)}$ denote the value of the matrix $\boldsymbol{D}$ in the VertexTestHat at line 9 after going through the loop starting on line 9 k times. Let $\boldsymbol{C}^{(k)}$ denote the matrix $\boldsymbol{C}$ in the VertexTest algorithm at line 16 after going through the loop starting on line 16 k times. Let $A_i^{(k)} = \{P_j : \boldsymbol{D}_{i,j}^{(k)} = 1\}$. Let $B_i^{(k)} = \{P_j : \boldsymbol{C}_{i,j}^{(k)} = 1\}$ as in the proof of Lemma 8.

Claim 1 of the proof of Lemma 8 shows that $Q_i \in B_i^{(k)}$ for all $k \in [L]$ and all $i \in [L]$. Therefore, it suffices to show that $A_i^{(k)} \supset B_i^{(k)}$ for all k and i. We prove this inductively. We claim that, after the first loop, $A_i^{(0)} \supset B_i^{(0)}$: by the inequality 4, if $\kappa^*(\tilde{P}_j \,|\, Q_i) = 0$, then for all $\epsilon > 0$, with probability tending 1 as $\boldsymbol{n} \longrightarrow \infty$, FaceContainHat$(\tilde{P}_j, Q_i \,|\epsilon)$ returns 0. Therefore, VertexTestHat sets an entry $\boldsymbol{D}_{i,j}^{(0)}$ to 0 in line 5 only if VertexTest sets the entry $\boldsymbol{C}_{i,j}^{(0)}$ to 0 in line 12. This proves the base case.

Now, suppose that $A_i^{(k)} \supset B_i^{(k)}$. Suppose $\boldsymbol{D}_{i,:}^{(k)} = \boldsymbol{e}_j^T$ for some j and i. There must be some j such that $\boldsymbol{C}_{i,j}^{(k)} = 1$ since $(Q_1, \ldots, Q_L)^T$ are a permutation of $(P_1, \ldots, P_L)^T$ by hypothesis and Claim 1 of the proof of 8. Then, by the inductive hypothesis, we must have $\boldsymbol{C}_{i,:}^{(k)} = \boldsymbol{e}_j^T$. Therefore, in the $(k+1)$th iteration, VertexTestHat sets an entry $\boldsymbol{D}_{i,j}^{(k)}$ to 0 in line 12 only if VertexTest sets the entry $\boldsymbol{C}_{i,j}^{(k)}$ to 0 in line 19. Hence, we delete an element P_l from $A_i^{(k)}$ in the $(k+1)$th iteration of VertexTestHat only if we delete P_l from $B_i^{(k)}$ in the $(k+1)$th iteration of VertexTest. This gives $A_i^{(k+1)} \supset B_i^{(k+1)}$. This proves the inductive step, completing the proof.

□

Theorem 5. *Let $\epsilon > 0$. Suppose that $P_1, \ldots, P_L$ satisfy **(A″)**, $\boldsymbol{\Pi}$ satisfies **(B′)**, and $\boldsymbol{S}$ satisfies **(C)** and **(D)**. Then, with probability tending to 1 as $\boldsymbol{n} \longrightarrow \infty$, PartialLabelHat$(\boldsymbol{S}, (\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)^T)$ returns $(\widehat{Q}_1, \ldots, \widehat{Q}_L)^T$ such that for every $i \in [L]$,*

$$\sup_{S \in \mathcal{S}} |\widehat{Q}_i(S) - P_i(S)| < \epsilon.$$

Theorem 5. Let $(\widehat{Q}_1, \ldots, \widehat{Q}_L) \longleftarrow$ DemixHat$(\tilde{P}_1^\dagger, \ldots, \tilde{P}_L^\dagger)$. By Theorem 4, there exists a permutation $\sigma : [L] \longrightarrow [L]$ such that for every $i \in [L]$,

$$\sup_{S \in \mathcal{S}} |\widehat{Q}_i(S) - P_{\sigma(i)}(S)| < \epsilon.$$

From the proof of Theorem 4, it is clear that each $\widehat{Q}_i$ is a ResidueHat estimator. The assumptions of Lemma 16 are satisfied. The result follows immediately from Lemma 16. □

References

S. Arora, R. Ge, and A. Moitra. Learning Topic Models–going beyond SVD. *Foundations of Computer Science*, 2012.

S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, and M. Zhu. A Practical Algorithm for Topic Modeling with Provable Guarantees. *Proceedings of the 30th International Conference on Machine Learning*, 2013.

S. Axler. *Linear Algebra Done Right*. Springer, 3 edition, 2015.

L. Berkman, B. H. Singer, and K. Manton. Black/White Differences in Health Status and Mortality among the Elderly. *Demography*, 26:661678, 1989.

G. Blanchard and C. Scott. Decontamination of Mutually Contaminated Models. *Proc. 17th Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, pages 1–9, 2014.

G. Blanchard, G. Lee, and C. Scott. Semi-Supervised Novelty Detection. *Journal of Machine Learning Research*, 11:2973–3009, 2010.

D. Blei, A. Ng, and M. Jordan. Latent Dirichlet Allocation. *Journal of Machine Learning research*, 3: 993–1022, 2003.

T. Cour, B. Sapp, and B. Taskar. Learning from partial labels. *Journal of Machine Learning*, 12:1501–1536, 2011.

W. Ding, M. Rohban, P. Ishwar, and V. Saligrama. Topic Discovery through Data Dependent and Random Projections. *Proceedings of the 30th International Conference on Machine Learning*, 2013.

W. Ding, M. Rohban, P. Ishwar, and V. Saligrama. Efficient Distributed Topic Modeling with Provable Guarantees. *In Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, pages 167–175, 2014.

D. Donoho and V. Stodden. When does Non-negative Matrix Factorization give a Correct Decomposition into Parts? *Advances in neural information processing systems*, 2003.

S. Jain, M. White, M. W. Trosset, and P. Radivojac. Nonparametric semi-supervised learning of class proportions. *arXiv preprint arXiv:1601.01944*, 2016.

R. Jin and Z. Ghahramani. Learning with multiple labels. *Advances in Neural Information Processing Systems*, pages 897–904, 2002.

F. Li and P. Perona. A Bayesian Hierarchical Model for Learning Natural Scene Categories. *Computer Vision and Pattern Recognition*, 2005.

L.-P. Liu and T. G. Dietterich. A conditional multinomial mixture model for superset label learning. *Advances in Neural Information Processing Systems*, pages 557–565, 2012.

L.-P. Liu and T. G. Dietterich. Learnability of the superset label learning problem. *Proceedings of the 31st International Conference on Machine Learning*, 32, 2014.

M.-T. Luong, R. Socher, and C. D. Manning. Better word representations with recursive neural networks for morphology. *CoNLL-2013*, 104, 2013.

N. Nyugen and R. Caruana. Classification with partial labels. *Proceedings of the 14th International Conference on Knowledge Discovery and Data Mining*, pages 551–559, 2008.

B. Poczos, L. Xiong, D. J. Sutherland, and J. Schneider. Nonparametric kernel estimators for image classification. *Computer Vision and Pattern Recognition (CVPR)*, 2012.

J. K. Pritchard, M. Stephens, N. A. Rosenberg, and P. Donnelly. Association Mapping in Structured Populations. *American Journal of Human Genetics*, 67:170181, 2000.

B. Recht, C. Re, J. Tropp, and V. Bittorf. Factoring Non-negative Matrices with Linear Programs. *In Advances in Neural Information Processing Systems*, pages 1214–1222, 2012.

T. Sanderson and C. Scott. Class Proportion Estimation with Application to Multiclass Anomaly Rejection. *Proc. 17th Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, W&CP 33:850–858, 2014.

C. Scott. A Rate of Convergence for Mixture Proportion Estimation, with Application to Learning from Noisy Labels. *Proc 18th Int. Conf. on Artificial Intelligence and Statistics*, pages 838–846, 2015.

C. Scott, G. Blanchard, and G. Handy. Classification with Asymmetric Label Noise: Consistency and Maximal Denoising. *Proc. Conf. on Learning Theory (COLT)*, 30:489–511, 2013.